Using Semidefinite Programming to Bound Distributions in Chemical Engineering Systems

by

Garrett Ryan Dowdy

Submitted to the Department of Chemical Engineering in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Chemical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2019

© Massachusetts Institute of Technology 2019. All rights reserved.

Author

Department of Chemical Engineering October 11th, 2018

Certified by..... Paul I. Barton Lammont du Pont Professor of Chemical Engineering Thesis Supervisor

Accepted by Patrick S. Doyle

Chairman, Department Committee on Graduate Theses

Using Semidefinite Programming to Bound Distributions in Chemical Engineering Systems

by

Garrett Ryan Dowdy

Submitted to the Department of Chemical Engineering on October 11th, 2018, in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Chemical Engineering

Abstract

Distributions appear in many forms in models of chemical engineering systems. Such distributions account for microscopic variability in the system while simultaneously explaining its macroscopic properties. These macroscopic properties are often of practical engineering interest. Thus, it is valuable to be able to characterize the underlying distributions that affect them.

Recently, in the mathematical programming literature, it was shown that it is possible to optimize a linear objective over a set of distributions by solving a specific type of convex optimization problem called a semidefinite program (SDP). From a theoretical perspective, SDPs can be solved efficiently. Furthermore, there exist several off-the-shelf codes designed specifically to solve SDPs.

This thesis demonstrates how these theoretical and practical advancements can be applied to chemical engineering problems featuring distributions. Broadly speaking, it shows how, given limited information about a distribution, one can use SDPs to calculate mathematically rigorous bounds on various descriptions of that distribution.

Two specific types of distributions are examined: particle size distributions and probability distributions arising in stochastic chemical kinetics, with the majority of the thesis covering the latter topic. The SDP-based bounding method described herein provides a rigorous solution to the long-standing "moment closure problem" arising in stochastic chemical kinetics. Moreover, it provides a means of analyzing of stochastic chemical kinetic systems which cannot be effectively analyzed using existing methods. The bounding method does have some limitations, and we present several refinements of the method aimed at overcoming these limitations. Finally, we discuss several ideas through which the bounding method may be further improved, which have not yet been explored.

Thesis Supervisor: Paul I. Barton Title: Lammont du Pont Professor of Chemical Engineering Baby steps... All I have to do is take one little step at a time, and I can do anything.

BOB WILEY

Acknowledgments

A lot of people helped me on the way to my PhD, sometimes directly but more often indirectly, by just being a source of happiness in my life. In particular, I would like to thank the following people, in (mostly) random order:

- Paul Barton, for taking a gamble on me, and accepting me as your student, nearly 5 years ago. You and I weren't always best friends, but I appreciate what you did for me as my research advisor. In particular, first, you made sure that I always had funding. I never had to write a proposal just to keep my project going. Second, every time I gave you a paper to read, you read it very carefully, thinking through every statement I made. I didn't appreciate just how much endurance that required until I had to proofread someone's mathematical writing myself. Third, you gave me a lot of freedom in pursuing an area that interested me, even though it was a little outside of the group's comfort zone. For all this, thank you.
- Rahul Misra, for being my running buddy for the past several years. 7:30 AM seemed earlier and earlier each year, didn't it? These opportunities to let off some steam (both with the run itself and the post-run venting sessions) helped keep me sane.
- Ananth Govind Rajan, for a being a friend with whom I felt I could talk about anything, with no fear of judgment, and no filter. I'm going to miss our walks around the Charles.
- Stuart Harwood, for being my "mathematical big brother" when I first joined the Barton group, and patiently answering my countably infinite questions.
- Omar Labban, for being my gardening buddy. That watermelon was delicious.
- Evan Waldheter, for coaching me through a particularly trying time at MIT.
- Jeffrey Heier, for sewing my retina back onto my eyeball, when it very inconveniently decided to peel off during finals week of the first semester.

- Lucas at the Koch Cafe, for your generous "half-sized" portions of food.
- Harry Watson, for introducing me to the idea that "perfect is the enemy of good".
- Jose Gomez, for providing dating advice and good-humored distractions from research. Thank you, also, for being my guide and translator during our trip to Spain, and for supporting my dancing habits.
- Rohit Kannan, for your wry sense of humor, and for being the first person I turned to with my tricky mathematical questions in the middle years of my PhD.
- David Qiu, for being a constant presence at Ashdown events, and for alerting me to opportunities for free food.
- Amir Akbari, for being my companion on many little weekend excursions.
- Matias Vikse, for keeping me entertained with a steady stream of outrageous statements, and for teaching me useful Norwegian expressions like "handsome in biscuits", "hey haircut", and "slippers hero".
- Leia Dwyer, for including me in her social circles, especially during the first year, and for introducing me to many of the lovely hiking trails that New Hampshire has to offer.
- Kristen Severson, for being a lunch buddy, a fellow yogi, and excellent host for many game nights and get-togethers at her place.
- Max Jordan, for being my sailing buddy.
- Zsigi Varga, for being my fellow 10.34 TA. You did a good job making sure that we met all the deadlines, while I was busy wringing my hands over mathematical details. I'll never forget that congratulatory cake.
- Maria Tou, for being a pleasant teammate at Practice School and travel companion in Barcelona.

- Fara Lyu, for being a pleasant teammate at Practice School, and the leader of my favorite project team.
- Francesco for showing courage in speaking openly about mental health, and making it easier for me to do the same.
- Shawn Ferrulo, for patiently answering my seemingly endless list of healthrelated questions.
- Johannes Jaeschke, for being a pleasant guest in 66-270, and for valuing your family above your work.
- Peter Stechlinksi, for memorable mathematical conversations on the whiteboard in the lab, and for graciously hosting me and Matias on a weekend trip to Maine.
- Alin Tomescu, for being one of my regular lunch buddies, and for being honest about how you feel, at the risk of not being politically correct.
- Andrew Zmolek, for being a good roommate, and for commiserating with me about grad student life.
- Andrew Fiore, for being an enthusiastic attendee of "Puzzle Lunch Tuesdays", and for providing the puzzling group with high-quality wax color print-outs for those puzzles for which the usual black and white simply wouldn't cut it.
- Brinda Monian, for introducing me to the simple pleasure of caring for houseplants.
- Lenny Ferrari, for shooting the breeze with me in Ashdown's courtyard about the progress of our respective gardens.
- Lucas Caretta, for being a lunch buddy and a reminder of good times in Baytown, Texas.
- Martin Copenhaver, for being a reminder of good times at Georgia Tech, and for pointing me in the right direction when I was searching for a job.

- Mike, for shooting the breeze with me in the hallways of Building 66.
- Natasha Seelam, for helping me get REFS-X up and running.
- Nick Schickel, for being one of the most reliable attendees of "Puzzle Lunch Tuesdays".
- Richard Zhang, for being a good roommate, and for tutoring me with impromptu math lessons on the whiteboard in the kitchen.
- Sihang He, for joining me for breakfast in Ashdown's courtyard, and for being my companion on many little adventures during the Summer of 2018 including a memorable trip to Walden Pond.
- Tay Yang, for commiserating with me about PhD student life, and for reminding me of fun times at when we were roommates at Georgia Tech.
- Andrea Carney, for all the pleasant conversation when I visited Ashdown's front desk.
- Yijin Wei, for being my cooking buddy, for getting me involved in volunteering with GSC.
- John Gallagher, for hanging 5 pieces of art on the wall of my apartment, and for replacing the window screen draw-chain approximately 37 times.
- My grandfather, Bobby Dowdy, for our semi-regular phone calls, which both entertained me and helped me remember that there is a wide world beyond MIT.
- My grandmother, Peggy Hobbs, for our semi-regular phone calls, in which you never failed to remind me that I am loved.
- Mom, Dad, and my sister Lindsey, for our weekly "Family Skype Time", which helped remind me that life goes on outside of MIT; for our memorable trips to Hawaii, Banff, the Hike Inn, Cape Cod, the White Mountains, and Idaho (to

see the eclipse); for giving me a place to go home to every Christmas, where I could be a kid again, dropping the pretense of being a young professional; and, most importantly, for loving me regardless of the number of papers I have published and the number of awards I can list on my résumé.

Contents

| 1 | Introduction | | |
|----------|--------------|---|----|
| | 1.1 | Overview | 27 |
| | 1.2 | Thesis Contributions and Structure | 28 |
| 2 | Bou | ands on Particle Size Distributions | 31 |
| | 2.1 | Introduction | 31 |
| | | 2.1.1 Descriptions of PSDs | 34 |
| | 2.2 | Restrictions on Moment Sequences | 36 |
| | | 2.2.1 Notation | 36 |
| | | 2.2.2 LMI Constraints on Moments | 37 |
| | 2.3 | Bounds on Various Descriptions of PSDs | 38 |
| | | 2.3.1 Number of particles in a size interval | 38 |
| | | 2.3.2 Bounds on PSD histograms | 46 |
| | | 2.3.3 Bounds on the CDF | 48 |
| | | 2.3.4 Bounds on $D\alpha$ values | 54 |
| | 2.4 | More Moments Lead to Tighter Bounds | 59 |
| | | 2.4.1 Example | 60 |
| | 2.5 | Extensions | 61 |
| | 2.6 | Conclusion | 63 |
| 3 | Bou | unds on Stochastic Chemical Kinetic Systems at Steady State | 65 |
| | 3.1 | Introduction | 66 |
| | 3.2 | Mathematical Background | 70 |

| | 3.2.1 | Mathematical Notation | 70 |
|-----|--------|--|-----|
| | 3.2.2 | Stochastic Chemical Kinetics Notation | 70 |
| | 3.2.3 | Invariants and Independent Species | 70 |
| | 3.2.4 | The Chemical Master Equation | 72 |
| | 3.2.5 | Moments in Stochastic Chemical Kinetics | 73 |
| | 3.2.6 | The Closure Problem | 75 |
| | 3.2.7 | Closure Scheme Approximations | 76 |
| 3.3 | The E | Bounding Method | 77 |
| | 3.3.1 | The Paradigm | 77 |
| | 3.3.2 | Necessary Steady-State Moment Conditions | 79 |
| | 3.3.3 | A Semidefinite Program | 81 |
| | 3.3.4 | Conservatism in the Bounds | 82 |
| | 3.3.5 | Scaling | 84 |
| 3.4 | Mean- | -Bounding Examples | 85 |
| | 3.4.1 | Michaelis-Menten | 85 |
| | 3.4.2 | Reversible Dimerization | 86 |
| | 3.4.3 | Schlögl | 89 |
| | 3.4.4 | A Larger Example | 90 |
| 3.5 | Variar | nce Bounds | 92 |
| | 3.5.1 | An Upper-Bounding SDP | 94 |
| | 3.5.2 | A Lower-Bounding SDP? | 96 |
| 3.6 | Variar | nce-Bounding Examples | 97 |
| | 3.6.1 | Michaelis-Menten | 97 |
| | 3.6.2 | Reversible Dimerization | 97 |
| | 3.6.3 | Schlögl | 100 |
| | 3.6.4 | A Larger Example | 100 |
| 3.7 | Bound | ls on Probability | 100 |
| | 3.7.1 | An Upper Bounding SDP | 102 |
| | 3.7.2 | A Lower Bounding SDP | 103 |
| | 3.7.3 | Bounds on a Histogram | 103 |

| | 3.8 | Proba | bility-Bounding Examples | 104 |
|---|------|---------|---|-----|
| | | 3.8.1 | Michaelis-Menten | 105 |
| | | 3.8.2 | Reversible Dimerization | 105 |
| | | 3.8.3 | Schlögl | 108 |
| | | 3.8.4 | Larger Example | 110 |
| | 3.9 | Other | Probability Bounds | 110 |
| | 3.10 | Size of | f the Semidefinite Programs | 111 |
| | 3.11 | Conclu | usion | 112 |
| | 3.12 | Imple | mentation Details | 114 |
| 1 | Dur | omio | Bounda on Stochastic Chamical Kinotic Systems Usin | G |
| 4 | Sor | idofin | ite Programming | 115 |
| | Jen. | Introd | | 115 |
| | 4.1 | Mothe | metical Background | 110 |
| | 4.2 | 4 9 1 | Mathematical Notation | 110 |
| | | 4.2.1 | Stachagtic Chamical Vinatica Notation | 110 |
| | | 4.2.2 | Stochastic Chemical Kinetics Notation | 110 |
| | | 4.2.3 | Invariants and Independent Species | 117 |
| | | 4.2.4 | A Reduced State Space | 118 |
| | | 4.2.5 | The Chemical Master Equation | 119 |
| | | 4.2.6 | Moments in Stochastic Chemical Kinetics | 120 |
| | | 4.2.7 | The Closure Problem | 121 |
| | 4.3 | Bound | ls on Dynamic Systems | 122 |
| | | 4.3.1 | The Paradigm | 122 |
| | | 4.3.2 | Necessary Moment Conditions | 124 |
| | | 4.3.3 | Necessary Conditions on the Generalized Moments $\ . \ . \ .$. | 127 |
| | | 4.3.4 | A Semidefinite Program | 129 |
| | | 4.3.5 | Irrelevance of Intermediate Time Points | 129 |
| | | 4.3.6 | Inspiration from Previous Work | 130 |
| | | 4.3.7 | Choosing the Values of ρ | 132 |
| | | 4.3.8 | Bounds on Higher-Order Moments | 133 |

| | | 4.3.9 | Bounds on the Variance | 133 |
|----------|------|---------|--|-----|
| | | 4.3.10 | Conservatism in the Bounds | 135 |
| | | 4.3.11 | Scaling | 136 |
| | | 4.3.12 | Practical Applications | 136 |
| | 4.4 | Toy E | xample | 137 |
| | | 4.4.1 | Mean and Variance Bounds | 137 |
| | | 4.4.2 | Increasing the value of $m \ldots \ldots \ldots \ldots \ldots \ldots \ldots$ | 138 |
| | | 4.4.3 | Using more values of ρ | 138 |
| | 4.5 | A Bit | More Complexity | 143 |
| | | 4.5.1 | Mean and Variance Bounds | 143 |
| | | 4.5.2 | Using more values of ρ | 143 |
| | | 4.5.3 | Sensitivity of the Values of ρ | 145 |
| | 4.6 | Compl | ex Eigenvalues | 146 |
| | 4.7 | Perfect | t Bounds in the Absence of the Closure Problem | 150 |
| | 4.8 | Uncert | ainty in the Initial State | 152 |
| | 4.9 | A Larg | ger Example | 154 |
| | 4.10 | Compa | arison with SSA | 156 |
| | 4.11 | An Op | en System | 157 |
| | 4.12 | Next S | Steps: Bounds on Probability | 157 |
| | 4.13 | Conclu | usion | 158 |
| | 4.14 | Impler | nentation Details | 160 |
| 5 | Imp | roved | Bounds on Stochastic Chemical Kinetic Systems through | h |
| | an A | Alterna | ative Representation of State Space | 161 |
| | 5.1 | Introd | uction | 161 |
| | 5.2 | Notati | on | 162 |
| | 5.3 | Previo | us Work | 163 |
| | | 5.3.1 | Invariants and Reduced Reachable States | 163 |
| | | 5.3.2 | A Reduced Probability Distribution | 165 |
| | | 5.3.3 | Moments of the Reduced Distribution | 165 |

| | | 5.3.4 | The Bounding Paradigm | 166 |
|---|-----|--------|--|-----|
| | | 5.3.5 | Necessary Conditions on the Moment Sequences | 166 |
| | | 5.3.6 | Semidefinite Programs | 169 |
| | | 5.3.7 | $\bar{\mathcal{X}}$ as an Outer-Approximation of $\hat{\mathcal{X}}$ | 172 |
| | 5.4 | A Sec | ond Type of Outer-Approximation | 172 |
| | | 5.4.1 | A Pathological Example | 172 |
| | | 5.4.2 | Implications for the Bounds | 174 |
| | | 5.4.3 | Generalization | 176 |
| | 5.5 | An Al | ternative State Representation | 177 |
| | | 5.5.1 | Redundancy in the State Representation | 177 |
| | 5.6 | Altern | native Formulation of the Bounding Method | 179 |
| | | 5.6.1 | A New Probability Distribution | 179 |
| | | 5.6.2 | Moments | 179 |
| | | 5.6.3 | The Bounding Paradigm | 180 |
| | | 5.6.4 | Necessary Conditions on the Moment Sequence | 181 |
| | | 5.6.5 | Semidefinite Programs | 185 |
| | 5.7 | Revisi | ting the Pathological Example | 186 |
| | | 5.7.1 | Steady-State Analysis | 187 |
| | | 5.7.2 | Transient Analysis | 187 |
| | 5.8 | Discus | ssion | 188 |
| | | 5.8.1 | Weakness of the Steady-State Bounds | 188 |
| | | 5.8.2 | Strictly Better Bounds Are Not Guaranteed | 189 |
| | 5.9 | Conclu | usion | 189 |
| 6 | An | Additi | ional Steady-State Necessary Condition | 191 |
| | 6.1 | Introd | luction | 191 |
| | 6.2 | Derivi | ng the Necessary Condition | 192 |
| | 6.3 | Using | the Necessary Conditions as Constraints | 193 |
| | | 6.3.1 | Redundancy | 194 |
| | 6.4 | When | Does it Help? | 195 |

| | | 6.4.1 | A Rank Condition | 195 |
|---|--------------------|---------|--|-----|
| | | 6.4.2 | An Invariant Polynomial | 195 |
| | 6.5 | Exam | ple | 196 |
| | 6.6 | Exten | sions | 197 |
| | 6.7 | Conclu | usion | 198 |
| - | т | | | |
| 1 | Imp | oroved | Numerical Performance Through Linear Programming A | p- |
| | proz | | ons | 201 |
| | (.1 7 .0 | Introd | uction | 201 |
| | 7.2 | Notati | | 201 |
| | | 7.2.1 | Invariants and Reduced Reachable States | 202 |
| | 7.3 | Nume | rical Stability through an LP Relaxation | 204 |
| | | 7.3.1 | Relaxation Yields Valid Bounds | 204 |
| | | 7.3.2 | Relaxing the LMIs | 205 |
| | | 7.3.3 | Choosing \mathcal{P} and \mathcal{N}_c | 207 |
| | | 7.3.4 | Small Example | 210 |
| | | 7.3.5 | Large Example | 212 |
| | | 7.3.6 | Summary | 213 |
| | 7.4 | Iterati | ve Refinement of LP Bounds | 214 |
| | | 7.4.1 | Calculation of Cutting Planes | 214 |
| | | 7.4.2 | A Cutting Plane Algorithm | 215 |
| | | 7.4.3 | Convergence | 217 |
| | | 7.4.4 | Small Example | 220 |
| | | 7.4.5 | Large Example | 221 |
| | | 7.4.6 | Summary | 222 |
| | 7.5 | Warm | -Starting \mathcal{P} and \mathcal{N}_c | 223 |
| | | 7.5.1 | Derivation of the Dual SDP | 223 |
| | | 7.5.2 | Sum of Squared Polynomials Interpretation | 228 |
| | | 7.5.3 | Dual Equivalent of the LP Relaxation of the Primal SDP | 229 |
| | | 7.5.4 | The Warm-Starting Algorithm | 233 |

| | | 7.5.5 | Example | 235 |
|--------------|-----|---------|---|-----|
| | 7.6 | Conclu | sion | 237 |
| 8 | Con | clusior | 1 | 239 |
| | 8.1 | Summa | ary | 239 |
| | 8.2 | Future | Directions | 239 |
| \mathbf{A} | Sup | plemer | ntal Material for Chapter 2 | 243 |
| | | A.0.1 | Implementation Details | 243 |
| | | A.0.2 | The Problem with LP-based Methods | 243 |
| | | A.0.3 | Proof of Claim 1 | 245 |
| | | A.0.4 | Condition (2.7) implies nonnegativity of the moments | 247 |
| | | A.0.5 | Equivalence of Problems (2.10) and (2.11) | 247 |
| | | A.0.6 | Sufficiency of the k th order LMIs $\ldots \ldots \ldots \ldots \ldots \ldots$ | 248 |
| | | A.0.7 | Derivation of lower bounding SDP | 249 |
| | | A.0.8 | Inconsistency of the "Dirac Delta Function" | 251 |
| | | A.0.9 | Derivation of Problem (2.24) | 252 |
| | | A.0.10 | Proof of the Monotonicity of the CDF Bounds | 253 |
| | | A.0.11 | Algorithm for calculating a lower bound on the CDF $\ . \ . \ .$ | 255 |
| | | A.0.12 | Algorithm for calculating a lower bound on $D\alpha$ | 256 |
| в | Sup | plemer | ntal Material for Chapter 3 | 257 |
| | | B.0.1 | A Moment Enumeration Scheme | 257 |
| | | B.0.2 | Definition of Matrices $\mathbf{M}^{0}_{n}(\boldsymbol{\mu}), \mathbf{M}^{0}_{n-1}(\boldsymbol{\mu}), \text{and} \mathbf{M}^{\mathbf{e}_{j}}_{n-1}(\boldsymbol{\mu}) \ldots .$ | 263 |
| | | B.0.3 | Derivation of LMIs | 264 |
| | | B.0.4 | Proof that LMIs (3.27) - (3.29) Imply (3.24) - (3.26) | 268 |
| | | B.0.5 | Size of the SDPs | 271 |
| | | B.0.6 | Scaling | 272 |
| | | B.0.7 | LMIs to Exclude Non-integer States | 276 |
| | | B.0.8 | Rate Constants for the Larger Example | 278 |
| | | B.0.9 | Relative Probabilities for the Schlögl System | 279 |

| | | B.0.10 Number of States in the Larger Example |
|---|-----|---|
| | | B.0.11 Derivation of SDP (3.39) |
| С | Sup | blemental Material for Chapter 4 293 |
| | C.1 | Introduction $\ldots \ldots 293$ |
| | C.2 | Complex Eigenvalues |
| | | C.2.1 Linear Equations |
| | | C.2.2 Additional Constraints |
| | | C.2.3 Bringing It All Together |
| | | C.2.4 Complex Conjugates |
| | | C.2.5 An Augmented SDP 301 |
| | C.3 | The Absence of the Closure Problem |
| | | C.3.1 Theoretical Reasoning |
| | C.4 | Bounds on Higher-Order Moments |
| | C.5 | Proof of the Monotonicity of the Bounds with Increasing $m \ldots \ldots 309$ |
| D | Sup | blemental Material for Chapter 5 315 |
| | D.1 | Introduction $\ldots \ldots 315$ |
| | D.2 | A Moment Enumeration Scheme |
| | D.3 | Definition of $\mathbf{M}_{n}^{0}(\mathbf{y})$, $\mathbf{M}_{n-1}^{0}(\mathbf{y})$, and $\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\mathbf{y})$ |
| | D.4 | Derivation of LMIs |
| | D.5 | Construction of the C matrix $\ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots 316$ |
| | | D.5.1 Linear Mapping Between Moments |
| | | D.5.2 Linear Mapping Between Time Derivatives |
| | | D.5.3 Generalization |
| | D.6 | Construction of the $\mathbf{A}^{\mathcal{E}}$ matrix |
| | | D.6.1 Computing the $a_{r,i}$ Coefficients $\ldots \ldots \ldots \ldots \ldots \ldots 323$ |
| | D.7 | Revisiting a Non-Pathological Example |

List of Figures

| 2-1 | Upper bounds on a PSD Histogram | 48 |
|------|---|-----|
| 2-2 | Cartoon example of the "stair-step" functions guaranteed to bound the | |
| | true CDF | 51 |
| 2-3 | CDF bounds with few test points | 53 |
| 2-4 | Increasing the number of test points results in much smoother bound- | |
| | ing functions. | 54 |
| 2-5 | Cartoon example of the graphical analysis used to obtain (approxi- | |
| | mate) bounds on $D\alpha$ | 57 |
| 2-6 | Graphical analysis for obtaining (approximate) bounds on D50 \ldots . | 59 |
| 2-7 | Pointwise CDF bounds, varying the number of moments specified | 61 |
| 3-1 | Bounds on the Michaelis-Menten system at steady state | 86 |
| 3-2 | Bounds on the Reversible Dimerization system at steady state $\ . \ . \ .$ | 87 |
| 3-3 | Bounds on the Schlögl system at steady state | 91 |
| 3-4 | A larger reaction system | 92 |
| 3-5 | Bounds on the larger reaction system in Figure 3-4 at steady state. $% \left({{{\rm{A}}_{{\rm{B}}}} \right)$. | 93 |
| 3-6 | Bounds on variances for the Michaelis-Menten system at steady state | 98 |
| 3-7 | Bounds on variances for the Reversible Dimerization system at steady | |
| | state | 99 |
| 3-8 | Bounds on variances for the Schlögl system at steady state \ldots . | 101 |
| 3-9 | Bounds on variances for the larger reaction system in Figure 3-4 at | |
| | steady state | 101 |
| 3-10 | Bounds on a histogram for the Michaelis-Menten System at steady state | 106 |

| 3-11 | Bounds on a histogram for the Reversible Dimerization System at | |
|------|--|-----|
| | steady state | 107 |
| 3-12 | Initial bounds on a histogram for the Schlögl System at steady state . | 108 |
| 3-13 | Bounds on a histogram for the Schlögl System at steady state | 109 |
| 3-14 | Bounds on a histogram for the larger reaction system in Figure 3-4 at | |
| | steady state | 110 |
| 4-1 | Initial bounds on System (4.51) | 139 |
| 4-2 | Increasing the parameter m improves the quality of the bounds | 140 |
| 4-3 | Using an enlarged set $\mathcal R$ improves the quality of the bounds | 141 |
| 4-4 | Further enlarging the set \mathcal{R} further improves the quality of the bounds. | 142 |
| 4-5 | Using an enlarged \mathcal{R} improves the bounds | 144 |
| 4-6 | Using an even larger set $\mathcal R$ even further improves the bounds | 145 |
| 4-7 | A choice of \mathcal{R} that is slightly off does not dramatically deteriorate the | |
| | bounds. | 147 |
| 4-8 | A very poor choice of $\mathcal R$ weakens the bounds. \ldots \ldots \ldots \ldots | 148 |
| 4-9 | Accounting for complex eigenvalues improves bounds | 150 |
| 4-10 | Perfect bounds in the absence of the closure problem | 151 |
| 4-11 | Bounds on the mean and variance for Reaction System (4.53) with an | |
| | uncertain initial state. | 153 |
| 4-12 | A larger reaction system | 155 |
| 4-13 | Bounds on the mean molecular counts of species A and H for the | |
| | reaction system shown in Figure 4-12 | 156 |
| 4-14 | Bounds on the mean molecular counts of species A and H for the | |
| | reaction system shown in Figure 4-12 | 158 |
| 5-1 | Base-case time-varying bounds. | 176 |
| 5-2 | Improvement in transient bounds | 188 |
| 7-1 | Bounds on the Michaelis-Menten system at steady state | 211 |
| 7-2 | A larger reaction system | 212 |

| 7-3 | A larger reaction system | 213 |
|-----|--|-----|
| 7-4 | Cutting plane bounds on the Michaelis-Menten system at steady state. | 221 |
| 7-5 | Cutting plane bounds on the larger reaction system | 222 |
| 7-6 | Bounds on $\langle P \rangle_{ss}$ generated through two different algorithms | 237 |
| | | |
| B-1 | Relative probabilities for the Schlögl System over a range of parameter | |
| | values | 282 |
| B-2 | Alternative representation of Figure 3-4 | 283 |
| | | |
| C-1 | Time-varying bounds on a third-order moment. \ldots \ldots \ldots \ldots | 308 |
| C-2 | Time-varying bounds on a fourth-order moment | 308 |
| | | |
| D-1 | Time-varying bounds based on reaction extents. | 327 |

List of Tables

| 3.1 | Sizes of the SDPs solved in the examples of this chapter | 113 |
|-----|--|-----|
| 4.1 | The set of reachable states \mathcal{X} of the system described in Section 4.5 | 152 |
| 7.1 | Microscopic rate constants for the autoregulatory gene network | 236 |
| B.1 | Association between single indices $(j \in \{1, 2, 3,\})$ and multi-indices | |
| | $(\mathbf{j} \in \mathbb{N}^{\hat{N}})$ | 260 |
| B.2 | Reaction rate constants for the Larger Example $\ldots \ldots \ldots \ldots$ | 278 |

Chapter 1

Introduction

1.1 Overview

Many models of chemical engineering systems feature a distribution – for example, a crystal size distribution, molecular energy distribution, cell age distribution, probability distribution, polymer chain length distribution, and so on [53]. Such distributions account for microscopic variability in the system while simultaneously explaining its macroscopic behavior. For example, the molecular energy distribution can be linked to the temperature of a system [36]. Similarly, the dispersity of a polymer chain length distribution can be correlated with its mechanical properties [44]. Such macroscopic properties are often of practical engineering interest. Thus, it is valuable to be able to characterize the underlying distributions that affect them.

Relatively recently, Lasserre [40] showed that it is possible to optimize over a set of distributions by solving a specific type of convex optimization problem called a semidefinite program (SDP). From a theoretical standpoint, SDPs can be solved to ϵ -optimality in polynomial time [71]. From a practical standpoint, there exist several off-the-shelf solvers designed to handle SDPs specifically [66, 2, 76, 22]. This situation has led prominent figures in the mathematical programming community to make statements such as "semidefinite programs can be solved very efficiently, both in theory and in practice" [71].

These theoretical and practical advances suggested that there may be something to

be gained in applying SDPs to chemical engineering problems featuring distributions. Broadly speaking, the goal of this thesis was to explore that possibility.

1.2 Thesis Contributions and Structure

This thesis demonstrates how, given limited information about a distribution, we can use SDPs to calculate mathematically rigorous bounds on various descriptions of that distribution. In particular, we consider two questions:

- Given only the moments of a particle size distribution, what other details of the distribution can we infer?
- Given a stochastic chemical kinetic system, large enough that obtaining the exact solution is intractable, can we bound the solution?

The answer to the first question constitutes a relatively small portion of the thesis. It is covered entirely in Chapter 2. Specifically, in this chapter, we motivate the analysis of particle size distributions. Then, we show how SDPs can be used to bound the number of particles in an arbitrary size interval. Generalizing this result, we show that they can also be used to bound the industrially relevant D10, D50, and D90 of the distribution.

The answer to the second question accounts for the rest of the thesis. In Chapter 3, we review the background of stochastic chemical kinetics, including the moment closure problem. We then show how SDPs can be used to bound the stationary distribution of a stochastic chemical kinetic system. Specifically, we show how to bound the mean molecular count of each species and the variance in this count. Furthermore, we show how to bound the probability that the molecular count is in a particular range.

In Chapter 4, we explain how all of the results from Chapter 3 can be generalized to time-varying distributions. We show how each of the bounds mentioned in the preceding paragraph can be computed at an arbitrary time t, not just in the limit as $t \to +\infty$. In Chapter 5, we identify a weakness of the bounding method described in Chapters 3 and 4, and we present an alternative formulation of the method that does not exhibit the same weakness. This alternative formulation is based on expressing the state of the system in terms of extents of reaction, rather than molecular counts.

In Chapter 6, we describe a simple refinement of the steady-state bounding method described in Chapter 3, which hugely improves the quality of the bounds produced in stochastic chemical kinetic systems with a specific structure.

Finally, in Chapter 7, we describe a strategy for dealing with the numerical issues identified in previous chapters. In particular, we show how a prototypical bounding SDP can be outer approximated by a linear program (LP). We show how this LP approximation can be iteratively refined and intelligently initialized using the results of an inaccurate optimal solution of the original SDP.

Each of these chapters corresponds to a paper that has either been published in a peer-reviewed journal or is in preparation for submission.

Chapter 2

Bounds on Particle Size Distributions

Many chemical engineering processes involve a population of particles with a distribution of sizes that changes over time. Because calculating the time evolution of the full particle size distribution (PSD) is computationally expensive, it is common to instead calculate the time evolution of only finitely many moments of the distribution. The problem with moments is that they provide only a summary description of the PSD. In particular, they do not contain enough information to answer industrially relevant questions such as: How many particles are there in the size range [a,b]? What is the shape of the distribution? What is its D10? While these questions cannot be answered exactly, in this chapter, we demonstrate that one can efficiently calculate rigorous bounds on the answers by solving semidefinite programs.

The contents of this chapter were published as a peer-reviewed paper [12] in Chemical Engineering Science.

2.1 Introduction

Many chemical engineering processes involve a population of particles with a distribution of sizes that changes over time. For example, crystallization, colloidal suspension formation, catalyst attrition, polymerization, and aerosol formation all fit this general framework [53]. In each of these processes, the particle size distribution (PSD) can have a large effect on macroscopic properties of engineering interest. For example, for pharmaceutical crystals, the PSD affects the ease with which the crystals can be filtered and compacted into tablets, thereby affecting the cost and processing time of the pharmaceutical product [74, 50]. Moreover, once the drug has been introduced to a patient's system, its rate of dissolution (i.e. bioavailability) depends strongly on the PSD. Thus, the PSD is tied to both the pharmaceutical's efficacy and safety [6].

Because of the importance of the PSD in these diverse chemical engineering applications, many researchers have considered the question of how to model a PSD's time evolution. A common starting point is to describe the PSD using a *number density function*, which changes over time according to a partial differential equation (PDE). One can then solve the PDE to obtain a complete description of the PSD at each point in time (in terms of the number density function). In some cases, the PDE has an analytical solution. In most cases, though, the complexity of the model requires a numerical solution. Many algorithms for numerical solution of the PDE are available. However, they are computationally expensive, and usually too slow for on-line applications.

For this reason, instead of modeling the entire PSD, it is common to model only finitely many *moments* of the PSD. The advantage of this "method of moments" is that we no longer have to solve a PDE; instead, we need only solve a system of ordinary differential equations (ODEs) [32, 54]. This certainly reduces the computational burden, but it comes at a cost: moments are only a summary description of the PSD, i.e., they do not contain enough information to reconstruct all of its details. This is because there are, in general, many PSDs corresponding to a given finite set of moment values [46]. Thus, given only finitely many moments of an unknown distribution, there is no clear answer to industrially relevant questions such as:

- How many particles have size in the range *a* to *b*?
- What is the D10 of the distribution¹?

¹For those unfamiliar with this description of a PSD, it will be explained shortly

• What is the qualitative shape of the distribution?

Faced with these questions, one might be tempted to apply one of the various methods available for constructing a number density function with a specified finite set of moments [11, 33, 8, 75]. With the resulting number density function, answering the above questions would be trivial. However, the problem with this strategy should be clear from the foregoing discussion: the calculated number density function describes just one of the many PSDs with the specified moments. Accordingly, it would provide just one of the many possible (valid) answers to each of the above questions, giving us a false sense of certainty in our knowledge of the distribution.

In this chapter, we take a more rigorous approach. Acknowledging that reconstructing a PSD from finitely many moments is an ill-posed inverse problem, we make no attempt to answer the above questions exactly. Instead, we calculate provable *bounds* on the answers. These bounds require only finitely many moments of the distribution and no a priori knowledge of its shape. In principle, the moments could be obtained through either simulation or experimental measurement. However, due to the difficulties in accurately measuring the moments of a PSD [11], we expect the methods of this chapter to be most valuable to the community concerned with simulating moments.

The idea of bounding the PSD is not new. McGraw [45] and Shmakov [61] describe methods for calculating bounds on PSDs by solving Linear Programs (LPs). These methods have the appeal of flexibility and relative simplicity. However, the bounds that they produce are not truly rigorous, as explained in Appendix A.

The bounding methods presented in this chapter attain true rigor by employing more sophisticated mathematics. In particular, instead of solving LPs, we will calculate bounds by solving Semidefinite Programs (SDPs) [71]. These SDPs are a natural application of results from the mathematical literature regarding moments of positive finite Borel measures (i.e., generalized distributions) [42]. While SDPs have been applied in chemical engineering in the context of optimal control [70], to the best of the authors' knowledge, their natural application to particle size distributions has, until now, gone unnoticed.

2.1.1 Descriptions of PSDs

The foregoing section referenced several descriptions of PSDs. We now define these descriptions more precisely, as they will be important in what follows.

Number Density Functions

In chemical engineering applications, a PSD is commonly described by a number density function, f. This f is defined such that the integral $\int_a^b f(x)dx$ gives the number of particles with size in the range a to b. Typically, this is a number per unit volume (or mass) of the processing system. However, for the sake of brevity, we will drop the qualifier "per unit volume...". We will always be referring to a number concentration rather than a total number, and this will be clear from the units given.

While the mathematical properties of the function f are rarely explicitly stated, it is always assumed (and reasonably so) that:

- 1. f(x) is undefined for all x < 0, because it doesn't make physical sense to talk about particles with negative size.
- 2. f is nonnegative on \mathbb{R}_+ , because you cannot have a negative count of particles of any size.
- 3. The integral $\int_0^{+\infty} x^j f(x) dx$ is well defined and finite for all $j \in \{0, 1, 2, ...\}$. We will see shortly that this implies that moments of f are well defined and finite.

We will use the symbol $M^{\infty}(\mathbb{R}_+)$ to denote the set of functions that satisfy the above three properties. Occasionally, we will want to focus on those $f \in M^{\infty}(\mathbb{R}_+)$ which are nonnegative only on a specific interval $[a, b] \subset \mathbb{R}_+$ and zero elsewhere. We will use the symbol $M^{\infty}([a, b])$ to denote this subset of $M^{\infty}(\mathbb{R}_+)$.

Cumulative Distribution Function

Closely related to the number density function is the *cumulative distribution function* (CDF). The CDF is a function F defined such that F(x) is the number of particles

with size less than or equal to x. When a PSD is described by a number density function f, the CDF is given by

$$F(x) = \int_0^x f(x') dx'.$$
 (2.1)

The function F is necessarily nondecreasing. This follows from its physical definition, but it can also be seen as a consequence of the fact that f is nonnegative.

$\mathbf{D}\alpha$ Values

A PSD can also be described in terms of its D10, D50, and D90 values, which are the 10th, 50th, and 90th percentiles of the distribution, respectively. There is some ambiguity in the literature as to whether these percentiles are on a number or mass basis. However, the mass basis convention seems to be more common, so this is the convention we will adopt. To illustrate what we mean by "mass basis", if we say that the D10 of a PSD is 30 μ m, this means that 10% of the mass of the PSD is attributable to particles with size less than or equal to 30 μ m. We generalize this idea to define D α for an arbitrary $\alpha \in [0, 100]$. In particular, if the D α of a PSD is s, then α % of the mass of the PSD is attributable to particles with size less than or equal to s.

 $D\alpha$ values are not just of academic interest. In the pharmaceutical industry, quality specifications are often expressed in terms of D10, D50, and D90 [19].

Moments

The *j*th moment of a PSD described by a number density function $f \in M^{\infty}(\mathbb{R}_+)$ is defined as

$$\mu_j \equiv \int_0^{+\infty} x^j f(x) dx, \quad j = 0, 1, 2, \dots$$
 (2.2)

As discussed in the introduction, moments are a very common output of PSD simulations. Moreover, some moments can be measured [73, 11]. For these reasons, moments are worthy of special attention. In the next section, we examine some of their intriguing mathematical properties.

2.2 Restrictions on Moment Sequences

We suspect it is intuitive to most readers that the moments of a PSD cannot be arbitrary real numbers. Indeed, just by inspecting Equation (2.2) and considering that f must be nonnegative on \mathbb{R}_+ , one can conclude that each moment of a PSD described by f must be nonnegative:

$$\mu_j \ge 0$$
, for all $j \in \{0, 1, 2, ...\}.$ (2.3)

In fact, there are far more subtle constraints that any PSD moment sequence must satisfy. It turns out that an understanding of these constraints enables us to use the moments to calculate bounds on other descriptions of the PSD. Thus, this section is devoted to developing that understanding.

2.2.1 Notation

To describe the constraints on the moments of a PSD, we need to first define some mathematical notation. Let \mathbb{N} denote the set of natural numbers, $\{0, 1, 2, ...\}$. Let $\tilde{\mu}$ represent an infinite moment sequence $(\mu_0, \mu_1, \mu_2, ...)$. Then, for any $n \in \mathbb{N}$, the *n*th order *Hankel matrices* of $\tilde{\mu}$ are defined, element-wise, as

$$\begin{aligned} \mathbf{H}_{n}(\tilde{\mu})(i,j) &\equiv \mu_{i+j-2}, \\ \mathbf{B}_{n}(\tilde{\mu})(i,j) &\equiv \mu_{i+j-1}, \quad \text{for all } i,j \in \{1,...,n+1\} \end{aligned}$$

$$\begin{aligned} \mathbf{C}_{n}(\tilde{\mu})(i,j) &\equiv \mu_{i+j}. \end{aligned}$$

$$\end{aligned}$$

$$\begin{aligned} (2.4)$$
If we write these matrices out in full, it becomes apparent that they are symmetric and have a cross-diagonal pattern. For example,

$$\mathbf{H}_{n}(\tilde{\mu}) = \begin{bmatrix} \mu_{0} & \mu_{1} & \mu_{2} & \dots & \mu_{n} \\ \mu_{1} & \mu_{2} & \mu_{3} & \dots & \mu_{n+1} \\ \mu_{2} & \mu_{3} & \mu_{4} & \dots & \mu_{n+2} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mu_{n} & \mu_{n+1} & \mu_{n+2} & \dots & \mu_{2n} \end{bmatrix} \in \mathbb{R}^{(n+1)\times(n+1)}.$$
(2.5)

For a symmetric matrix such as $\mathbf{H}_n(\tilde{\mu})$, the notation $\mathbf{H}_n(\tilde{\mu}) \succeq \mathbf{0}$ means that $\mathbf{H}_n(\tilde{\mu})$ is positive semidefinite, or all its eigenvalues are nonnegative. This is equivalent to the statement

$$\mathbf{x}^{\mathrm{T}}\mathbf{H}_{n}(\tilde{\mu})\mathbf{x} \ge 0, \text{ for all } \mathbf{x} \in \mathbb{R}^{n+1}.$$
 (2.6)

The statement $\mathbf{H}_n(\tilde{\mu}) \succeq \mathbf{0}$ is a simple example of a *linear matrix inequality (LMI)* [71].

2.2.2 LMI Constraints on Moments

With these definitions, we can now state the first central claim of this chapter, which is closely related to Theorem 3.2 appearing in [42, Chapter 3] and Theorem 1.3 of [62]:

Claim 1. If $\tilde{\mu} = (\mu_0, \mu_1, \mu_2, ...)$ is a moment sequence for a PSD, then $\tilde{\mu}$ must satisfy

$$\mathbf{H}_n(\tilde{\mu}) \succeq \mathbf{0} \text{ and } \mathbf{B}_n(\tilde{\mu}) \succeq \mathbf{0}, \text{ for all } n \in \mathbb{N}.$$
 (2.7)

The proof of this claim is surprisingly simple and can be found in the Appendix.

What does Condition (2.7) mean? In short, it implies a large number of constraints on the sequence $\tilde{\mu}$. In particular, one can show that Condition (2.7) implies the nonnegativity of each μ_j , which we already deduced (see Appendix). The remaining constraints are complicated polynomials in $\tilde{\mu}$ and are much less intuitive. This complexity is intriguing, but irrelevant for our present purposes. All that matters at present is that the constraints can be represented in the concise LMI form in which they appear in Condition (2.7), and that this form is recognized by optimization codes [28, 43].

Claim 1 is concerned with PSDs supported on (i.e., spread over) all of \mathbb{R}_+ . However, as we will see shortly, we will also be interested in distributions confined to some subset $[a, b] \subset \mathbb{R}_+$, where a and b are bounds on the minimum and maximum particle size, respectively. As with the distributions supported on all of \mathbb{R}_+ , we also have a necessary condition for the moments of confined PSDs:

Claim 2. Let $0 \le a \le b$. If $\tilde{\mu} = (\mu_0, \mu_1, \mu_2, ...)$ is a moment sequence for a PSD where a and b are bounds on the minimum and maximum particle size, respectively, then $\tilde{\mu}$ must satisfy

$$\mathbf{H}_n(\tilde{\mu}) \succeq \mathbf{0} \quad and \quad (a+b)\mathbf{B}_n(\tilde{\mu}) - \mathbf{C}_n(\tilde{\mu}) - ab\mathbf{H}_n(\tilde{\mu}) \succeq \mathbf{0}, \text{ for all } n \in \mathbb{N}.$$
 (2.8)

The proof of this claim is similar to that of Claim 1.

2.3 Bounds on Various Descriptions of PSDs

We will now demonstrate how the LMI constraints on moment sequences described in the previous section can be used to bound various numerical descriptions of a PSD. In each of the sections below, we will assume that the only available description of the PSD is its first m + 1 moments $\mu_0, ..., \mu_m$, presumably obtained through simulation.

2.3.1 Number of particles in a size interval

Suppose that we have a PSD for which we know only the moments $\mu_0, ..., \mu_m$, and that we would like to calculate the number of particles with size in the range a to b, inclusive. If we knew the number density function $\hat{f} \in M^{\infty}(\mathbb{R}_+)$ describing the PSD, we could readily obtain the desired quantity via the integral $\int_a^b \hat{f}(x) dx$. However, given that we don't know \hat{f} , the problem is more complicated. As suggested in the introduction, the complication arises from the fact that there may be multiple number density functions $f \in M^{\infty}(\mathbb{R}_+)$ with the same moments $\mu_0, ..., \mu_m$ but differing values for the integral $\int_a^b f(x) dx$. For this reason, the problem of calculating the number of particles with size in the range [a, b] given the moments $\mu_0, ..., \mu_m$ is ill-posed.

The situation is not hopeless, though. Let N represent the number of particles in the range [a, b]. That is, let $N \equiv \int_a^b \hat{f}(x) dx$. While we cannot precisely calculate N, we will see that we can calculate *bounds* on N. That is, we can calculate two numbers \underline{N}^L and \overline{N}^U such that $\underline{N}^L \leq N \leq \overline{N}^U$, guaranteed.

The problem of calculating these bounds is closely related, mathematically, to the problem of calculating bounds on the probability that a real-valued random variable X, with specified moments $\mu_0, ..., \mu_m$, realizes a value in a given interval [a, b]. Lasserre [42, 41] provides an elegant solution to this problem, demonstrating that the probability bounds can be obtained by solving semidefinite programs (SDPs). Lasserre's arguments are highly technical, being based on the mathematically sophisticated concept of a positive finite Borel measure. The following section relies heavily on ideas found in [42, 41]. However, in an effort to make these ideas more accessible to a general chemical engineering audience, we avoid discussion of measures, presenting the arguments instead in terms of more familiar number density functions.

The Upper Bound SDP

Let us first focus our attention on the upper bound, \overline{N}^U . Let N^U be the maximum² value of $\int_a^b f(x)dx$ obtained in considering all number density functions f with the moments $\mu_0, ..., \mu_m$. This definition can be expressed formally as

$$N^{U} \equiv \max_{f} \quad \int_{a}^{b} f(x)dx$$

s.t. $f \in M^{\infty}(\mathbb{R}_{+}),$
 $\int_{0}^{+\infty} x^{j}f(x)dx = \mu_{j}, \quad \forall j \in \{0, ..., m\}.$ (2.9)

²Technically, the supremum of the values of $\int_a^b f(x)dx$. However, we omit this technicality, assuming that most readers will be more familiar with the term "maximum".

The decision variable of this optimization problem is a function f rather than a finite-dimensional vector.

Since the unknown number density function \hat{f} is an element of $M^{\infty}(\mathbb{R}_+)$ and $\int_0^{+\infty} \hat{f}(x) dx = \mu_j$ for all $j \in \{0, ..., m\}$, it follows that \hat{f} is feasible for Problem (2.9). Thus, N^U is an upper bound on $\int_a^b \hat{f}(x) dx = N$, as desired. If we could calculate N^U , we would be done. However, it is not obvious how to solve Problem (2.9).

Suppose that we decompose each $f \in M^{\infty}(\mathbb{R}_+)$ into f = g + h, where both g and h are nonnegative number density functions. We will require that g(x) = 0 outside of [a, b]. In other words, g is confined to [a, b], in the sense described in the previous section. We will require no such confinement for h. Formally, we have $g \in M^{\infty}([a, b])$ and $h \in M^{\infty}(\mathbb{R}_+)$. We can then rewrite Problem (2.9) equivalently as

$$N^{U} = \max_{g,h} \quad \int_{a}^{b} g(x)dx + \int_{a}^{b} h(x)dx$$

s.t. $g \in M^{\infty}([a,b]), \quad h \in M^{\infty}(\mathbb{R}_{+}),$
 $\int_{0}^{+\infty} x^{j}g(x)dx + \int_{0}^{+\infty} x^{j}h(x)dx = \mu_{j}, \quad \forall j \in \{0,...,m\}.$ (2.10)

Furthermore, one can show that the above problem is equivalent to

$$N^{U} = \max_{g,h} \int_{0}^{+\infty} g(x)dx$$

s.t. $g \in M^{\infty}([a,b]), h \in M^{\infty}(\mathbb{R}_{+}),$
 $\int_{0}^{+\infty} x^{j}g(x)dx + \int_{0}^{+\infty} x^{j}h(x)dx = \mu_{j}, \forall j \in \{0,...,m\},$ (2.11)

where the only change is in the objective function. (For a proof of this equivalence, see the Appendix.)

Now, let $\tilde{y} = (y_0, y_1, y_2, ...)$ and $\tilde{z} = (z_0, z_1, z_2, ...)$ be the moment sequences for

the functions h and g, respectively. We can then write

$$N^{U} = \max_{g,h,\tilde{y},\tilde{z}} \quad z_{0}$$

s.t. $g \in M^{\infty}([a,b]), \quad h \in M^{\infty}(\mathbb{R}_{+}),$
 $z_{j} + y_{j} = \mu_{j}, \quad \forall j \in \{0,...,m\},$
 $\int_{0}^{+\infty} x^{j}g(x)dx = z_{j}, \quad \int_{0}^{+\infty} x^{j}h(x)dx = y_{j}, \quad \forall j \in \mathbb{N}.$ (2.12)

By Claims 1 and 2, we can conclude that \tilde{y} and \tilde{z} must necessarily satisfy Conditions (2.7) and (2.8), respectively. We can thus add in these redundant conditions without changing the feasible set of the optimization problem:

$$N^{U} = \max_{g,h,\tilde{y},\tilde{z}} \quad z_{0}$$

s.t. $g \in M^{\infty}([a,b]), \quad h \in M^{\infty}(\mathbb{R}_{+}),$
 $z_{j} + y_{j} = \mu_{j}, \quad \forall j \in \{0,...,m\},$
 $\int_{0}^{+\infty} x^{j}g(x)dx = z_{j}, \quad \int_{0}^{+\infty} x^{j}h(x)dx = y_{j}, \quad \forall j \in \mathbb{N},$
 $\mathbf{H}_{n}(\tilde{y}) \succeq \mathbf{0}, \quad \mathbf{B}_{n}(\tilde{y}) \succeq \mathbf{0}, \quad \forall n \in \mathbb{N},$
 $\mathbf{H}_{n}(\tilde{z}) \succeq \mathbf{0}, \quad (a+b)\mathbf{B}_{n}(\tilde{z}) - \mathbf{C}_{n}(\tilde{z}) - ab\mathbf{H}_{n}(\tilde{z}) \succeq \mathbf{0}, \quad \forall n \in \mathbb{N}.$
(2.13)

The reader may naturally wonder where all this mathematical manipulation is headed. After all, comparing Problems (2.13) and (2.9), it may seem that we have only managed to complicate things. In fact, we have prepared ourselves for a crucial step: the difficulty of Problem (2.9) is attributable to the decision variables g and h; we now remove these variables from Problem (2.13), obtaining a simpler problem:

$$\overline{N}^{U} \equiv \max_{\tilde{y}, \tilde{z}} \quad z_{0}$$
s.t. $z_{j} + y_{j} = \mu_{j}, \quad \forall j \in \{0, ..., m\},$

$$\mathbf{H}_{n}(\tilde{y}) \succeq \mathbf{0}, \quad \mathbf{B}_{n}(\tilde{y}) \succeq \mathbf{0}, \quad \forall n \in \mathbb{N},$$

$$\mathbf{H}_{n}(\tilde{z}) \succeq \mathbf{0}, \quad (a+b)\mathbf{B}_{n}(\tilde{z}) - \mathbf{C}_{n}(\tilde{z}) - ab\mathbf{H}_{n}(\tilde{z}) \succeq \mathbf{0}, \quad \forall n \in \mathbb{N}.$$
(2.14)

In Problem (2.14), we are no longer requiring that the moment sequences \tilde{y} and \tilde{z}

correspond to number density functions. Of course, the moment sequences corresponding to number density functions are still feasible, but the feasible set now also includes moment sequences corresponding to distributions which *cannot* be described by number density functions. For some readers, the idea of a distribution that cannot be described by a number density function may be foreign. We will see an example of such a distribution shortly. At present, though, all that matters is that we have enlarged the set of feasible moment sequences. This means that Problem (2.14) is a relaxation of Problem (2.13), and its optimal value is potentially greater. In terms of the newly introduced symbol \overline{N}^U , we have $\overline{N}^U \geq N^U$. Thus, \overline{N}^U serves as an upper bound on the number of particles in the interval [a, b]. It is perhaps not the tightest upper bound, as a result of including those distributions which cannot be described by number density functions. However, when we come to the examples, we will see that the conservatism introduced by including these distributions is reasonable.

Problem (2.14) is not only visually simpler than Problem (2.13); it is also very nearly an SDP. It differs from an SDP in two respects: first, the decision variables \tilde{y} and \tilde{z} are both infinite sequences; second, the last two lines of the problem statement describe an infinite number of LMI constraints. Both of these difficulties can be alleviated by truncating the sequences of LMI constraints. For example, instead of requiring that $\mathbf{H}_n(\tilde{y}) \succeq \mathbf{0}$ for all $n \in \mathbb{N}$, we will only require that $\mathbf{H}_k(\tilde{y}) \succeq \mathbf{0}$ for some specific $k \in \mathbb{N}$. One can show that this implies $\mathbf{H}_n(\tilde{y}) \succeq \mathbf{0}$ for all $n \leq k$ (see Appendix). The other three sequences of LMIs can be truncated similarly. We then need only consider as many moments $y_0, ..., y_r$ and $z_0, ..., z_r$ as are needed to fill out the resulting finite-dimensional Hankel matrices. In other words, the decision variables become the truncated sequences $y = (y_0, ..., y_r)$ and $z = (z_0, ..., z_r)$ for some finite $r \in \mathbb{N}$.

Truncating the LMIs in this way results in an SDP. In principle, this SDP is a relaxation of Problem (2.14), because we have relaxed constraints describing the feasible set. However, if we are careful about how we truncate the LMIs, the resulting SDP gives the same optimal value as Problem (2.14) (See Theorems 7.2(b) and 3.2 of [42]). In particular, if we let $k = \lceil \frac{m}{2} \rceil$, and r = 2k, then \overline{N}^U is given by the following

SDP:

$$\overline{N}^{U} = \max_{y,z} \quad z_{0}$$
s.t. $z_{j} + y_{j} = \mu_{j}, \quad \forall j \in \{0, ..., m\},$

$$\mathbf{H}_{k}(y) \succeq \mathbf{0}, \quad \mathbf{B}_{k-1}(y) \succeq \mathbf{0},$$

$$\mathbf{H}_{k}(z) \succeq \mathbf{0}, \quad (a+b)\mathbf{B}_{k-1}(z) - \mathbf{C}_{k-1}(z) - ab\mathbf{H}_{k-1}(z) \succeq \mathbf{0}.$$
(2.15)

Reviewing the inequalities we have established, we see

$$N \le N^U \le \overline{N}^U. \tag{2.16}$$

Thus, \overline{N}^U is an upper bound on N, as desired. Moreover, because \overline{N}^U is the optimal value of an SDP, we can calculate it efficiently using optimization codes such as CVX [28], using SeDuMi [66] as the underlying solver.

The Lower Bound SDP

Analogous to Problem (2.15), there is an SDP for calculating a *lower* bound on N:

$$\underline{N}^{L} \equiv \min_{z,y,w} \qquad z_{0}$$
s.t. $z_{j} + y_{j} + w_{j} = \mu_{j}, \quad \forall j \in \{0, ..., m\},$

$$\mathbf{H}_{k}(z) \succeq \mathbf{0}, \quad \mathbf{B}_{k-1}(z) \succeq \mathbf{0}, \qquad (2.17)$$

$$\mathbf{H}_{k}(y) \succeq \mathbf{0}, \quad a\mathbf{B}_{k-1}(y) - \mathbf{C}_{k-1}(y) \succeq \mathbf{0},$$

$$\mathbf{H}_{k}(w) \succeq \mathbf{0}, \quad \mathbf{B}_{k-1}(w) - b\mathbf{H}_{k-1}(w) \succeq \mathbf{0}.$$

The derivation of this SDP is similar to that of (2.15) (see Appendix). The primary difference is that three number density functions and three corresponding moment sequences (z, y, w) are required.

Scaling

As we will see in the coming examples, it is common for the moments of a PSD to range over many orders of magnitude. This can lead to numerical difficulties if one attempts to solve SDPs such as Problems (2.15) and (2.17) directly. Fortunately, there is an approach for scaling the moments that alleviates these numerical difficulties: define the scaled moments as $\mu'_j \equiv (\mu_j/\mu_0)c^j$ for $j \in \{0, ..., m\}$ and $c = \mu_0/\mu_1$. As described in [34], this amounts to normalizing the PSD and then scaling the size domain. The result is that the scaled moments are all approximately equal one (in fact, $\mu'_0 = \mu'_1 = 1$ exactly). Of course, since we've scaled the size domain, we must also scale the limits of the interval [a, b]. In particular, we let $a' \equiv ca$ and $b' \equiv cb$. This scaling of the problem data results in a scaled SDP. For example, the scaled version of Problem (2.15) is

$$\frac{\overline{N}^{U}}{\mu_{0}} = \max_{y',z'} \quad z'_{0}$$
s.t. $z'_{j} + y'_{j} = \mu'_{j}, \quad \forall j \in \{0, ..., m\},$

$$\mathbf{H}_{k}(y') \succeq \mathbf{0}, \quad \mathbf{B}_{k-1}(y') \succeq \mathbf{0},$$

$$\mathbf{H}_{k}(z') \succeq \mathbf{0}, \quad (a'+b')\mathbf{B}_{k-1}(z') - \mathbf{C}_{k-1}(z') - a'b'\mathbf{H}_{k-1}(z') \succeq \mathbf{0}.$$
(2.18)

Note that the optimal values has been divided by μ_0 to account for the normalization of the PSD.

We will not discuss this scaling further, as it is just a matter of numerics. However, in the examples that follow, whenever we speak of solving an SDP, it should be understood that we mean the appropriately *scaled* SDP.

Example

The authors of [17] use the method of moments in combination with computational fluid dynamics to describe the agglomeration and breakage of catalyst particles in a fluidized bed reactor. In Figure 5 of that paper, they report the time-varying moments of the catalyst PSD for one of their simulations. At the final time point, the moments

$$\mu_0 = 6.73 \times 10^3 \text{ cm}^{-3},$$

$$\mu_1 = 2.28 \times 10^2 \text{ cm}^{-2},$$

$$\mu_2 = 9.15 \times 10^0 \text{ cm}^{-1},$$

$$\mu_3 = 3.85 \times 10^{-1}.$$

(2.19)

Suppose we are interested in the number of particles with size in the range 85 μ m to 150 μ m, which we will denote by N. By solving Problem (2.15), we can compute an upper bound on this quantity: $N \leq \overline{N}^U = 1.85 \times 10^3 \text{ cm}^{-3}$.

One might wonder about the quality of this bound. That is, one might wonder if the calculated \overline{N}^U is a gross overestimation of what is possible given the moments in Equation (2.19), or if there does indeed exist a PSD which satisfies (2.19) and which has \overline{N}^U particles in the range 85 μ m to 150 μ m. In fact, there is such a distribution. Consider the PSD in which there are exactly three different sizes represented:

$$x_1 = 85 \ \mu m,$$

 $x_2 = 150 \ \mu m,$ (2.20)
 $x_3 = 428.04685945926 \ \mu m,$

with corresponding concentrations

$$a_{1} = 5.46816164144708 \times 10^{2} \text{ cm}^{-3},$$

$$a_{2} = 1.308027852880654 \times 10^{3} \text{ cm}^{-3},$$

$$a_{3} = 4.875155982974638 \times 10^{3} \text{ cm}^{-3}.$$
(2.21)

One can verify that this PSD has the moments given by (2.19) (error $< 1 \times 10^{-9}$ %). Moreover, the number of particles with size between 85 µm to 150 µm is simply $a_1 + a_2 = 1.85 \times 10^3$ cm⁻³. Comparing this value with \overline{N}^U , we see that the calculated upper bound is indeed tight; there is no overestimation.

In developing Problem (2.15), we allowed for distributions which cannot be de-

are:

scribed by a number density function and pointed out that doing so might introduce reasonable conservatism into the bound \overline{N}^U . The distribution of the previous paragraph nicely illustrates this point. Technically speaking, this discrete distribution cannot be described by a number density function³ Such an infinitely sharp bimodal distribution is unlikely to be observed in any real industrial processing system. Thus, considering the discrete distribution as a feasible solution introduces some conservatism into the bound on the number of particles in the interval. It is, in some sense, a "worst-case" bound. However, given that the worst-case discrete distribution is physically well-defined (see the description in the previous paragraph) and matches the specified moments (the only specification), we believe that this conservatism is reasonable.

We can also use Problem (2.17) to compute a lower bound on N. In this case, the lower bound of $\underline{N}^{L} = 1.12 \times 10^{-5} \text{ cm}^{-3} \approx 0 \text{ cm}^{-3}$ is trivial. The physical interpretation, roughly speaking, is that there exists a PSD with the specified moments and with no particles in the specified size range. In our experience, this will be the case whenever the size interval under consideration is small relative the overall spread of the distribution. If we expand our consideration to all particles in the interval 85 μ m to 450 μ m, we can solve Problem (2.17) to compute a nontrivial lower bound: $\underline{N}^{L} = 1.77 \times 10^{3} \text{ cm}^{-3}$.

2.3.2 Bounds on PSD histograms

The ability to calculate upper and lower bounds on the number of particles in an arbitrary interval [a, b] has some interesting implications. First, it allows us to generate bounds on a PSD histogram. In principle, one could choose a sequence of intervals $[a_i, b_i]$, corresponding to the bins of a histogram, and calculate the bounds \underline{N}_i^L and \overline{N}_i^U for each bin. In practice, the lower bounds will usually be trivial, since, by design, each histogram bin covers only a small fraction of the total spread of the distribution. Thus, the lower bounds are omitted in Algorithm 1, which formalizes the idea.

 $^{^{3}}$ One could attempt to describe the distribution in terms of "Dirac Delta functions". However, the concept of a Dirac Delta *function* is mathematically inconsistent (see the Appendix for details).

Algorithm 1 Generating a PSD Histogram

Input:

- 1. Moments $\mu_0, ..., \mu_m$.
- 2. Estimate of the maximum particle size, x_{max} (e.g. $4\mu_1/\mu_0$).
- 3. Number of histogram bins, n.

Output: An upper bound \overline{N}_i^U on the number of particles in each bin.

Algorithm: Set $\Delta x := x_{\max}/n$. for i = 1, ..., n do Set $a_i := (i - 1)\Delta x$. Set $b_i := i\Delta x$. Solve Problem (2.15) to calculate \overline{N}_i^U . end for

Example

The authors of [60] model a hen egg-white lysozyme crystallization process, describing the PSD using only its first five moments. Figure 6 of that paper provides the data resulting from one of their moment simulations, and we can use this data to illustrate the utility of Algorithm 1. Focusing on the final time point of the "true" (black) curve, we can read off the moments of the final PSD:

$$\mu_{0} = 1.1 \times 10^{3} \text{ mL}^{-1},$$

$$\mu_{1} = 1.8 \times 10^{1} \text{ cm/mL},$$

$$\mu_{2} = 3.3 \times 10^{-1} \text{ cm}^{2}/\text{mL},$$

$$\mu_{3} = 6.4 \times 10^{-3} \text{ cm}^{3}/\text{mL},$$

$$\mu_{4} = 1.3 \times 10^{-4} \text{ cm}^{4}/\text{mL}.$$
(2.22)

If we apply Algorithm 1 to the moments given in (2.22), using $x_{\text{max}} = 0.05$ cm and n = 10. we get the graphical output shown in Figure 2-1. This output gives us some idea of the shape of the underlying distribution, which is not apparent from the moments alone.



Figure 2-1: Upper bounds on each bin of a histogram for any PSD with moments given by (2.22).

One must be careful about the interpretation of results such as those in Figure 2-1. The upper bounding bars are not, themselves, a histogram. Rather, the values of \overline{N}_i^U place bounds on the height of each bar of a histogram for any PSD with moments given by (2.22).

2.3.3 Bounds on the CDF

The ability to calculate bounds on the number of particles in an arbitrary size interval also allows us to calculate pointwise bounds on the cumulative distribution function (CDF). Recall that the CDF evaluated at x is nothing more than the number of particles with size less than or equal to x. In other words, it is the number of particles in the interval [0, x]. We can thus calculate bounds on CDF(x) by solving Problems (2.15) and (2.17) for the special case where a = 0 and b = x. In this case, Problem (2.15) reduces to

$$\overline{N}^{U}(x) \equiv \max_{y,z} \quad z_{0}$$
s.t. $z_{j} + y_{j} = \mu_{j}, \quad \forall j \in \{0, ..., m\},$

$$\mathbf{H}_{k}(y) \succeq \mathbf{0}, \quad \mathbf{B}_{k-1}(y) \succeq \mathbf{0},$$

$$\mathbf{H}_{k}(z) \succeq \mathbf{0}, \quad x\mathbf{B}_{k-1}(z) - \mathbf{C}_{k-1}(z) \succeq \mathbf{0}.$$
(2.23)

Moreover, while less obvious, it can be shown (see Appendix) that Problem (2.17) reduces to

$$\underline{N}^{L}(x) \equiv \min_{z,w} \quad z_{0}$$
s.t. $z_{j} + w_{j} = \mu_{j}, \quad \forall j \in \{0, ..., m\},$

$$\mathbf{H}_{k}(z) \succeq \mathbf{0}, \quad \mathbf{B}_{k-1}(z) \succeq \mathbf{0},$$

$$\mathbf{H}_{k}(w) \succeq \mathbf{0}, \quad \mathbf{B}_{k-1}(w) - x\mathbf{H}_{k-1}(w) \succeq \mathbf{0}.$$
(2.24)

Suppose we solve Problems (2.23) and (2.24) for each x_i in a set $\{x_i\}_{i=1}^n$ such that each $x_i \ge 0$. We then have upper an lower bounds on the CDF at each point x_i . This in itself is interesting information. However, it turns out that we can use this information to compute something more interesting still. Since the CDF function is necessarily a nondecreasing function, we can interpolate (in some sense) between the points x_i to construct functions CDF_n^U and CDF_n^L which are guaranteed to bound the CDF function at every x in the interval $[0, \max_i x_i]$. Moreover, we can refine our bounding functions as much as we like by introducing additional x_i and calculating the corresponding bounds on $\text{CDF}(x_i)$.

To see this, suppose, given a finite set of moments $\mu_0, ..., \mu_m$, we solve Problems (2.23) and (2.24) for some $x_1 \ge 0$, obtaining $\overline{N}^U(x_1)$ and $\underline{N}^L(x_1)$, upper and lower bounds, respectively, on $\text{CDF}(x_1)$. Since the CDF is a nondecreasing function, we can then immediately conclude that $\text{CDF}(x) \le \overline{N}^U(x_1)$ for all $x \in [0, x_1]$. By similar reasoning, we can also conclude that $\text{CDF}(x) \ge \underline{N}^L(x_1)$ for all $x \in [x_1, +\infty)$. We can then construct two functions CDF_1^U and CDF_1^L which bound the true CDF function at every x in the finite interval $[0, x_1]$. In particular, we let

$$CDF_{1}^{L}(x) \equiv \begin{cases} 0, & \text{if } x \in [0, x_{1}), \\ \underline{N}^{L}(x_{1}), & \text{if } x = x_{1}, \end{cases}$$
(2.25)

and

$$CDF_1^U(x) \equiv \left\{ \overline{N}^U(x_1), \quad \text{if } x \in [0, x_1]. \right.$$
(2.26)

Suppose we then evaluate $\overline{N}^U(x_2)$ and $\underline{N}^L(x_2)$ for some $x_2 \in [0, x_1)$. By the same reasoning we applied to x_1 , we can conclude that $\text{CDF}(x) \leq \overline{N}^U(x_2)$ for all $x \in [0, x_2]$ and that $\text{CDF}(x) \geq \underline{N}^L(x_2)$ for all $x \in [x_2, +\infty)$. Moreover, we are guaranteed that $\overline{N}^U(x_2) \leq \overline{N}^U(x_1)$ and $\underline{N}^L(x_2) \leq \underline{N}^L(x_1)$ (see Appendix). With this information, we can refine our bounding functions. In particular, we let

$$CDF_{2}^{L}(x) \equiv \begin{cases} 0, & \text{if } x \in [0, x_{2}), \\ \underline{N}^{L}(x_{2}), & \text{if } x \in [x_{2}, x_{1}), \\ \underline{N}^{L}(x_{1}), & \text{if } x = x_{1}, \end{cases}$$
(2.27)

and

$$\operatorname{CDF}_{2}^{U}(x) \equiv \begin{cases} \overline{N}^{U}(x_{2}), & \text{if } x \in [0, x_{2}], \\ \overline{N}^{U}(x_{1}), & \text{if } x \in (x_{2}, x_{1}]. \end{cases}$$
(2.28)

These "stair-step" functions are nondecreasing, like the unknown CDF they bound. Graphically, they would look something like the cartoon functions shown in Figure 2-2.



Figure 2-2: Cartoon example of the "stair-step" functions guaranteed to bound the true CDF.

One can imagine generating bounds $\overline{N}^U(x_i)$ and $\underline{N}^L(x_i)$ for additional test points x_i in the interval $[0, x_1]$. Each additional test point x_i would subdivide one of the current large steps into two smaller steps, thereby producing a smoother bounding function.

This idea of constructing bounding functions can be formalized as an algorithm. In particular, Algorithm 2 describes a procedure for constructing the upper bounding function, CDF_n^U . The algorithm automatically selects a sequence of test points $\mathbf{x} = (x_0, x_1, ..., x_n)$, with the selection criterion given in the last line of the "for" loop. Geometrically, each test point is selected as the midpoint of the current largest step⁴ ; this helps to produce a relatively smooth bounding function. The procedure for constructing the lower bounding function, CDF_n^L , is similar, and can be found in the Appendix (Algorithm 6).

Example

The authors of [7] model an ice cream crystallization process, describing the PSD using only its first four moments. Figure 6 of that paper provides the data resulting from one of their moment simulations, and we can use this data to illustrate the

⁴In the algorithm, we use the symbol " \in " rather than ":=" to account for the possibility that there may be multiple "largest steps". This is unlikely to actually occur in practice. However, if it does, the midpoint of any of these largest steps may be selected as the next midpoint.

Algorithm 2 Generating an upper bound on the CDF

Input:

- 1. Moments $\mu_0, ..., \mu_m$.
- 2. Estimate of the maximum particle size, x_{max} .
- 3. Number of test points, n.

Output: A function CDF_n^U which provides an upper bound for the true CDF on the interval $[0, x_{\text{max}}]$.

Algorithm:

Set $x_1 := 0$. Set $u_1 := x_1$. Solve Problem (2.23) to calculate $\overline{N}^U(x_1)$. Set $x_2 := x_{\max}$. Set $u_2 := x_2$. Solve Problem (2.23) to calculate $\overline{N}^U(x_2)$. Set $j^* := 1$. for i = 3, ..., n do Set $x_i := \frac{1}{2}(u_{j^*} + u_{j^*+1})$. Solve Problem (2.23) to calculate $\overline{N}^U(x_i)$. Set $\mathbf{u} := \operatorname{sort}(\mathbf{x})$. Set $j^* \in \arg \max_{j \le i-1} \left\{ \left(\overline{N}^U(u_{j+1}) - \overline{N}^U(u_j) \right) (u_{j+1} - u_j) \right\}$ end for Set $\operatorname{CDF}_n^U(x) := \begin{cases} \overline{N}^U(u_1), & \text{if } x = 0, \\ \overline{N}^U(u_j), & \text{if } x \in (u_{j-1}, u_j]. \end{cases}$ (2.29)

utility of Algorithms 2 and 6. Focusing on an arbitrarily selected time point of 750 s, we can read off the following moment values:

$$\mu_0 = 7.2 \times 10^{14} \text{ m}^{-3},$$

$$\mu_1 = 6.1 \times 10^9 \text{ m/m}^3,$$

$$\mu_2 = 5.9 \times 10^4 \text{ m}^2/\text{m}^3,$$

$$\mu_3 = 6.8 \times 10^{-1} \text{ m}^3/\text{m}^3.$$

(2.30)

If we apply Algorithms 2 and 6 to the moments given in (2.37), using $x_{\text{max}} = 4 \times 10^{-5}$ m and n = 10, we get the results shown in Figure 2-3. If we increase the number of test points to n = 30, we get the smoother curves shown in Figure 2-4. We wish to emphasize that the true CDF is guaranteed to lie between the two bounding

functions. Moreover, these two functions provide an excellent qualitative impression of the shape of the PSD.



Figure 2-3: Bounding functions for the CDF of a PSD with moments given in (2.19). These bounding functions were generated using n = 10 test points, and are thus relatively rough.



Figure 2-4: Increasing the number of test points to n = 30 results in much smoother bounding functions. Compare with Figure 2-3.

2.3.4 Bounds on $D\alpha$ values

Lastly, the ability to calculate bounds on the number of particles in an arbitrary interval can be extended to allow us to calculate bounds on D10, D50, and D90 values.

Recall that $D\alpha = s$ is equivalent to the statement that α percent of the mass of the PSD is attributable to particles with size less than or equal to s. Assuming that \hat{f} is a number density function describing the PSD in terms of a characteristic length, and that each particle can be described entirely by this characteristic length, the foregoing statement can be expressed mathematically as

$$\frac{\int_0^s x^3 \hat{f}(x) dx}{\int_0^{+\infty} x^3 \hat{f}(x) dx} = \frac{\alpha}{100}.$$
(2.31)

Now, if we don't know the the true number density function \hat{f} describing the PSD, and we know only the moments $\mu_0, ..., \mu_m$, we are interested in the following question: what is the minimum/maximum value of s for which there exists a number density function f which satisfies equation (2.31) and which has moments $\mu_0, ..., \mu_m$? This question can be formalized as an optimization problem. For example, the minimization problem is

$$D\alpha^{L}(\alpha) \equiv \min_{s,f} \quad s$$

s.t. $f \in M^{\infty}(\mathbb{R}_{+}), \quad s \in \mathbb{R}_{+},$
$$\frac{\int_{0}^{s} x^{3} \hat{f}(x) dx}{\int_{0}^{+\infty} x^{3} \hat{f}(x) dx} = \frac{\alpha}{100},$$

$$\int_{0}^{+\infty} x^{j} f(x) dx = \mu_{j}, \quad \forall j \in \{0, ..., m\}.$$

$$(2.32)$$

It turns out that Problem (2.32) is difficult to solve directly, but we can make progress if we invert the problem. Instead of asking "what is the minimum value of s associated with α ?" we will ask "what is the maximum value of α associated with s?". Then, the relevant optimization problem is

$$\alpha^{U}(s) \equiv \max_{f} \quad 100 \frac{\int_{0}^{s} x^{3} \hat{f}(x) dx}{\int_{0}^{+\infty} x^{3} \hat{f}(x) dx}$$

s.t. $f \in M^{\infty}(\mathbb{R}_{+}),$
 $\int_{0}^{+\infty} x^{j} f(x) dx = \mu_{j}, \quad \forall j \in \{0, ..., m\}.$ (2.33)

Problem (2.33) is more tractable than Problem (2.32) essentially because it more closely resembles the problems we've encountered before. In particular, it is similar to Problem (2.9), with the only difference appearing in the objective function. Because of this similarity, we manipulate Problem (2.33) in much the same way that we did Problem (2.9): decomposing f into two functions g and h, with g confined to the interval [0, s]; adding in the necessary conditions for the moments of g and h; relaxing the problem by removing the references to g and h altogether; and finally truncating the LMI constraints. This manipulation results in the following SDP:

$$\overline{\alpha}^{U}(s) \equiv \max_{z,y} \quad (100/\mu_{3})z_{3}$$
s.t. $z_{j} + y_{j} = \mu_{j}, \quad \forall j \in \{0, ..., m\},$

$$\mathbf{H}_{k}(y) \succeq \mathbf{0}, \quad \mathbf{B}_{k-1}(y) \succeq \mathbf{0},$$

$$\mathbf{H}_{k}(z) \succeq \mathbf{0}, \quad s\mathbf{B}_{k-1}(z) - \mathbf{C}_{k-1}(z) \succeq \mathbf{0}.$$
(2.34)

As before, we let $k = \lceil \frac{m}{2} \rceil$.

By similar reasoning, one can derive the following SDP for calculating a lower bound on the α associated with a given s:

$$\underline{\alpha}^{L}(s) \equiv \min_{z,y} \quad (100/\mu_{3})y_{3}$$

s.t. $z_{j} + y_{j} = \mu_{j}, \quad \forall j \in \{0, ..., m\},$
 $\mathbf{H}_{k}(y) \succeq \mathbf{0}, \quad \mathbf{B}_{k-1}(y) \succeq \mathbf{0},$
 $\mathbf{H}_{k}(z) \succeq \mathbf{0}, \quad \mathbf{B}_{k-1}(z) - s\mathbf{H}_{k-1}(z) \succeq \mathbf{0}.$ (2.35)

These problems are interesting, but what about the original problem of calculating bounds on $D\alpha$ for specific values of α ? We are actually much closer to the solution of this problem than one might suspect.

To see this, consider what happens if we solve Problems (2.34) and (2.35) for a range of s values and plot the resulting $\overline{\alpha}^U(s)$ and $\underline{\alpha}^L(s)$ values. We will get two curves, $\overline{\alpha}^U$ and $\underline{\alpha}^L$. Roughly speaking, to obtain the desired bounds on $D\alpha$, all one has to do is draw a horizontal line across the $\overline{\alpha}^U$, $\underline{\alpha}^L$ plot so that it intersects the vertical axis at α . The value of s at which that line intersects the $\overline{\alpha}^U$ curve is $\underline{D\alpha}^L$, a lower bound on $D\alpha$; similarly, the value of s at which that line intersects the $\underline{\alpha}^L$ curve is $\overline{D\alpha}^U$, an upper bound on $D\alpha$. This idea is demonstrated in Figure 2-5.



Figure 2-5: Cartoon example of the graphical analysis used to obtain (approximate) bounds on $D\alpha$.

This graphical analysis may suffice for some purposes. However, if precision is necessary, we can take advantage of the fact that $\overline{\alpha}^U$ and $\underline{\alpha}^L$ are nondecreasing functions and calculate ϵ -close outer approximations of $\underline{D}\underline{\alpha}^L$ and $\overline{D}\overline{\alpha}^U$, through an efficient bisection procedure. The bisection procedure for $\overline{D}\overline{\alpha}^U$ is shown in Algorithm 3; the analogous procedure for $\underline{D}\underline{\alpha}^L$ can be found in the Appendix (Algorithm 7).

Example

In [52] paracetamol is crystallized from methanol and water, and the progress of the crystallization is monitored by focused beam reflectance measurement (FBRM). Specifically, the chord length distribution (CLD) is measured with FBRM, and this CLD is used to back-calculate the first five moments of the PSD. Figure 1 of that paper provides the data resulting from one such experiment, and we can use this data to illustrate the utility of Algorithm 3. Focusing on the final time point, we can read

Algorithm 3 Generating an upper bound on $D\alpha^U$

Input:

- 1. Moments $\mu_0, ..., \mu_m$.
- 2. The value of $\alpha \in [0, 100]$ for which you want to calculate bounds on $D\alpha$.
- 3. A tolerance $\epsilon > 0$.

Output: $\overline{D\alpha}^U_{\epsilon}$, which is an ϵ -close upper bound on $\overline{D\alpha}^U$. That is, it satisfies

$$\overline{\mathrm{D}\alpha}^{U}_{\epsilon} - \epsilon \leq \overline{\mathrm{D}\alpha}^{U} \leq \overline{\mathrm{D}\alpha}^{U}_{\epsilon}.$$
(2.36)

•

Algorithm:

```
Set \underline{D}\alpha_{\epsilon}^{U} := 0.

Set s := \mu_{1}/\mu_{0}.

Solve Problem (2.35) to obtain \underline{\alpha}^{L}(s).

while \underline{\alpha}^{L}(s) < \alpha do

Set \underline{D}\alpha_{\epsilon}^{U} := s.

Set s := 2s.

Solve Problem (2.35) to obtain \underline{\alpha}^{L}(s).

end while

Set \overline{D}\alpha_{\epsilon}^{U} := s.

while \overline{D}\alpha_{\epsilon}^{U} - \underline{D}\alpha_{\epsilon}^{U} > \epsilon do

Set s := (\overline{D}\alpha_{\epsilon}^{U} + \underline{D}\alpha_{\epsilon}^{U})/2.

Solve Problem (2.35) to obtain \underline{\alpha}^{L}(s).

if \underline{\alpha}^{L}(s) \geq \alpha then

Set \overline{D}\alpha^{U} := s.

else

Set \underline{D}\alpha^{U} := s.

end if

end while
```

off the following moments of the PSD:

$$\mu_{0} = 7.2 \times 10^{6} \text{ kg methanol}^{-1},$$

$$\mu_{1} = 1.2 \times 10^{3} \text{ m/kg methanol},$$

$$\mu_{2} = 3.4 \times 10^{-1} \text{ m}^{2}/\text{kg methanol},$$

$$\mu_{3} = 1.2 \times 10^{-4} \text{ m}^{3}/\text{kg methanol},$$

$$\mu_{4} = 5.0 \times 10^{-8} \text{ m}^{4}/\text{kg methanol}.$$
(2.37)

Solving Problems (2.34) and (2.35) with the moments given in (2.37) for a range of s values, we obtain the $\overline{\alpha}^U$ and $\underline{\alpha}^L$ curves shown in Figure 2-6. Drawing a horizontal line at $\alpha = 50$, and dropping verticals where it crosses the $\overline{\alpha}^U$ and $\underline{\alpha}^L$ curves, we see that $\overline{D50}^U \approx 5.8 \times 10^{-4}$ m and $\underline{D50}^L \approx 2.6 \times 10^{-4}$ m.

Applying Algorithm 3 to the moments given in (2.37), with $\alpha = 50$ and $\epsilon = 1 \times 10^{-6}$, we obtain $\overline{\text{D50}}_{\epsilon}^{U} = 5.76 \times 10^{-4}$ m. Note that this result is in agreement with the graphical analysis demonstrated in Figure 2-6.



Figure 2-6: Graphical analysis for obtaining (approximate) bounds on D50 for any PSD with moments given by (2.37)

2.4 More Moments Lead to Tighter Bounds

The more moment values that are specified, the less uncertainty there is in the underlying PSD. Thus, we would expect that knowing more moments of a PSD would allow us to compute tighter bounds on its various descriptions. This is indeed the case, as we demonstrate with the following example.

2.4.1 Example

Suppose we have a PSD described by a log-normal number density function

$$f(x) = \frac{N}{x\sigma\sqrt{2\pi}} \exp\left[-\frac{(\ln x - M)^2}{2\sigma^2}\right],$$
(2.38)

where N = 1000 particles per mL, $\sigma = 0.25$, M = 0, and the particle size x has units of μ m. The *j*th moment for this distribution is given by

$$\mu_j = N \exp\left[jM + \frac{1}{2}j^2\sigma^2\right]$$
(2.39)

With this expression, we can generate a partial moment sequence $\mu_0, ..., \mu_m$ for any desired value of m, and then apply Algorithms 2 and 6 to calculate pointwise bounds on the CDF of the distribution. By using different values of m, we can see how the quality of the bounds depends on the number of specified moments.

The results of such an experiment are shown in Figure 2-7. As expected, the quality of the bounds improves as we increase the value of m. The most dramatic improvement occurs in transitioning from m = 1 to m = 2. Thereafter, the improvements are more subtle. Of course, the bounds continue to improve, but the benefit of specifying each additional moment declines rapidly. In particular, the bounds computed using m = 17 (not shown) were not appreciably better than those computed using m = 6. With m = 18, we began to encounter numerical difficulties.



Figure 2-7: Pointwise CDF bounds calculated using Algorithms 2 and 6 with a varying number of moments specified. On each plot, the blue curve is the lower bound, the green curve is the upper bound, and the red curve is the analytical CDF.

2.5 Extensions

In this chapter, we have assumed that the moments of the PSD are known exactly. However, it may be the case that there is some uncertainty in the moment values. This uncertainty will likely be expressed in terms of a confidence interval, so that, for each $j \in \{0, ..., m\}$, we have $\mu_j^L \leq \mu_j \leq \mu_j^U$ with C% confidence.

It turns out that the method that we have described can handle this situation quite nicely. In each of the SDPs we have described, one simply replaces the equality constraints with the corresponding inequality. For example, Problem (2.15) becomes

$$\overline{N}^{U} = \max_{y,z} \quad z_{0}$$
s.t.
$$\mu_{j}^{L} \leq z_{j} + y_{j} \leq \mu_{j}^{U}, \quad \forall j \in \{0, ..., m\},$$

$$\mathbf{H}_{k}(y) \succeq \mathbf{0}, \quad \mathbf{B}_{k-1}(y) \succeq \mathbf{0},$$

$$\mathbf{H}_{k}(z) \succeq \mathbf{0}, \quad (a+b)\mathbf{B}_{k-1}(z) - \mathbf{C}_{k-1}(z) - ab\mathbf{H}_{k-1}(z) \succeq \mathbf{0}.$$
(2.40)

In fact, the method can also readily handle the more general situation where we have confidence *ellipsoids* for the moments, so that $\mu^{\mathrm{T}}\mathbf{V}^{-1}\mu \leq r$ with C% confidence, for some positive definite covariance matrix \mathbf{V} and positive scalar r. One cannot simply substitute $\mu^{\mathrm{T}}\mathbf{V}^{-1}\mu \leq r$ into the SDP, as the expression on the left-hand side of the inequality is nonlinear with respect to μ . However, by the Schur Complement Lemma [5], $\mu^{\mathrm{T}}\mathbf{V}^{-1}\mu \leq r$ is equivalent to

$$\begin{bmatrix} \mathbf{V} & \boldsymbol{\mu} \\ \boldsymbol{\mu}^{\mathrm{T}} & \boldsymbol{r} \end{bmatrix} \succeq \mathbf{0}, \tag{2.41}$$

which is a linear matrix inequality. Thus, we can substitute it into the SDPs we have described, replacing the equality constraints. For example, Problem (2.15) becomes

$$\overline{N}^{U} = \max_{y,z} \quad z_{0}$$
s.t.
$$\begin{bmatrix} \mathbf{V} & y+z \\ (y+z)^{\mathrm{T}} & r \end{bmatrix} \succeq \mathbf{0},$$

$$\mathbf{H}_{k}(y) \succeq \mathbf{0}, \quad \mathbf{B}_{k-1}(y) \succeq \mathbf{0},$$

$$\mathbf{H}_{k}(z) \succeq \mathbf{0}, \quad (a+b)\mathbf{B}_{k-1}(z) - \mathbf{C}_{k-1}(z) - ab\mathbf{H}_{k-1}(z) \succeq \mathbf{0}.$$
(2.42)

In making these substitutions, one must be careful about the interpretation of the resulting bounds. We can no longer say that the number of particles with size between a and b is less than or equal to \overline{N}^U , guaranteed; we must qualify this statement by

adding "with C% confidence".

2.6 Conclusion

In this chapter, we have demonstrated that semidefinite programming has a natural application to chemical engineering processes involving particle size distributions. In particular, given finitely many moments of a PSD, one can use semidefinite programming to obtain guaranteed bounds on several industrially relevant descriptions of the PSD, including:

- the number of particles in an arbitrary size interval,
- a particle size histogram,
- the cumulative distribution function,
- $D\alpha$ values in particular, D10, D50, and D90 values.

We have formulated SDPs and provided algorithms for calculating bounds on each of these descriptions. Because semidefinite programming is a special case of convex optimization, these bounds can be calculated efficiently (in tenths of seconds to seconds) using freely available optimization codes.

As suggested in the introduction, the problem of calculating bounds on a PSD is closely related to the well-studied problem of constructing a number density function with a specified finite set of moments. Many methods exist for solving this problem. Our own method does not replace these methods, but rather nicely complements them. While the existing methods focus attention on a single PSD which satisfies a set of moment constraints and is, in some sense, likely to occur, our method conservatively considers all possible PSDs with the specified moments. Together, the two types of methods give a more complete picture of the uncertainty surrounding the PSD.

The Matlab code used in the examples of this chapter is freely available for noncommercial uses as "PSD Bounding Tools" at https://yoric.mit.edu/software. Running this code requires that you have installed CVX, which is also freely available at http://cvxr.com/cvx/download.

Chapter 3

Bounds on Stochastic Chemical Kinetic Systems at Steady State

The method of moments has been proposed as a potential means to reduce the dimensionality of the chemical master equation (CME) appearing in stochastic chemical kinetics. However, attempts to apply the method of moments to the CME usually result in the so-called closure problem. Several authors have proposed moment closure schemes, which allow them to obtain approximations of quantities of interest, such as the mean molecular count for each species. However, these approximations have the dissatisfying feature that they come with no error bounds. This chapter presents a fundamentally different approach to the closure problem in stochastic chemical kinetics. Instead of making an approximation to compute a single number for the quantity of interest, we calculate mathematically rigorous bounds on this quantity by solving semidefinite programs (SDPs). These bounds provide a check on the validity of the moment closure approximations, and are in some cases so tight that they effectively provide the desired quantity. In this chapter, the bounded quantities of interest are the mean molecular count for each species, the variance in this count, and the probability that the count lies in an arbitrary interval. This chapter considers only steady-state probability distributions. The dynamic problem is discussed in Chapter 4.

The contents of this chapter were published as a peer-reviewed paper [14] in the

Journal of Chemical Physics and as a part of the conference proceedings [13] of 27th European Symposium on Computer Aided Process Engineering.

3.1 Introduction

Reacting chemical systems are traditionally modeled using deterministic rate equations. These equations implicitly assume that the state of the system is fully characterized by the concentration (i.e., molecular count) of each species, and that future concentrations are determined by the present concentrations. Of course, this is an approximation, because, at the microscopic scale, reactions require molecular collisions, and the frequency of these reactive collisions depends not only on the number of molecules of each species, but also on their positions and velocities [26]. Accordingly, if we predict future molecular counts of a reaction system based solely on present molecular counts, ignoring the finer details of molecular positions and velocities, we may find that the actual behavior of the physical system deviates from our prediction. For systems in which the number of reacting molecules is large, the variations resulting from the microscopic details of molecular positions and velocities tend to average out, so that the macroscopic system behaves in a predictable way. This accounts for the success and widespread use of deterministic rate equations in describing the time evolution of macroscopic reactive systems. However, when the number of reacting molecules is small (i.e., on the order of 10s or 100s), deterministic rate equations often fail to describe the system adequately, because the variations resulting from the neglected microscopic details do not average out; they become a significant and observable part of the system dynamics [47].

One way to deal with this situation is to model the microscopic details of the system explicitly, expressing the state of the system in terms of the position and velocity of every molecule, and then propagating the state through time by numerically integrating Newton's equations of motion. This is the paradigm of Molecular Dynamics. While it can be highly accurate, it also usually computationally expensive and slow [31]. Another approach is to express the state of the system in terms of only the molecular counts, and to account for the microscopic variation by modeling the system as a *stochastic process*, where the state of the system at one instant in time does not fully determine the state in the next instant [31]. This is the aim of Stochastic Chemical Kinetics, which, more precisely, models reacting systems as continuous time discrete state Markov processes.

In the Stochastic Chemical Kinetics paradigm, we assign a probability to each reachable state of the system. Collectively, these probabilities define a probability distribution. The way this distribution changes over time is elegantly described by the Chemical Master Equation (CME). If we could solve this CME, we would have a complete probabilistic description of the reaction system through time. In principle, this solution should be easy to obtain as the CME is conceptually simple: it is a linear time-invariant ordinary differential equation system. Yet, it is often quite difficult to solve for the reason that it contains an extraordinarily large number of reachable states and their corresponding unknown probabilities – even infinitely many [31].

The classical solution to this problem is to sample the reaction system using Gillespie's Stochastic Simulation Algorithm (SSA). While this algorithm is intuitively appealing and very easy to implement, it is often too slow in practice [27]. This is particularly true when the reaction system contains both fast and slow reactions. Many variants of Gillespie's algorithm have been developed with the aim of increasing its speed. Most of these involve some approximation that renders their results inexact and potentially misleading. Those that retain the exactness of Gillespie's algorithm remain fundamentally limited in that they must simulate every reaction [9].

Many other methods for approximately solving the CME have been proposed, which aren't based on Gillespie's SSA. For example, the chemical Langevin equation approach describes the time-evolution of the system using a stochastic differential equation [37, 48]. The Finite State Projection (FSP) method [49] and its variants [30, 72, 35] model only a finite subset of the reachable states explicitly, bounding the error that is introduced in neglecting the remaining states. The system size expansion approach splits the molecular counts into separate deterministic and fluctuating parts [69]. An overview of these and other methods is given by Schnoerr [59] and Higham [31]. While these methods effectively reduce the computational burden required to describe the stochastic chemical kinetic system, they are only approximations of the CME model.

Yet another method for dealing with the large number of states appearing in the CME is to model only finitely many moments of the time-varying probability distribution. In principle, the moments could then be used to calculate quantities of interest such as the mean count of molecules of each chemical species and the variances in these counts. The difficulty with this approach is that the system of ordinary differential equations describing how the moments change over time usually suffers from the so-called "closure problem", in which the time evolution of the moments up to order m depends on the values of moments up to order m + 1.

To deal with this problem, various authors have proposed "closure scheme" approximations. For example, C. S. Gillespie[25] assumes that the unknown distribution is normal; Smadbeck and Kaznessis[64] assume it is maximum entropy. With such assumptions, the low-order moments determine the entire distribution (including the high-order moments), and the closure problem is bypassed. One can then approximate both dynamic trajectories and steady-state values for the moments. Unfortunately, these are just approximations, with no bounds on error¹. The approximation could be good; on the other hand, it could be quite bad. These two scenarios are indistinguishable, unless some second, independent method is applied to the problem.

Recently, several authors [13, 56, 39, 24] (including ourselves) independently discovered and proposed an alternative strategy: instead of making an approximation to compute a single number for the quantity of interest, let us calculate mathematically rigorous bounds on the quantity. These bounds are computed by solving semidefinite programs (SDPs), where the decision variables are moments of the unknown distribution. This idea of moment-based SDPs was popularized by Lasserre [42] in the

¹One notable exception is the closure scheme described by Naghnaeian and Del Vecchio[51] which can provide error bounds under the condition that the molecular count of each species present in the system is bounded. However, the scalability of this method is doubtful from a theoretical perspective, as it requires solving a linear program (LP) whose size is proportional to the number of reachable states.

mathematical literature. Its application to stochastic chemical kinetics is appealing for three reasons: first, it makes no unprovable assumptions about the unknown distribution; second, it naturally fits the mathematical structure inherent in stochastic chemical kinetics (e.g., polynomial propensity functions and linearity of the CME); third, SDPs are convex optimization problems – meaning that, at least in theory, they can be solved efficiently [71].

This chapter is a more-detailed, tutorial version of our original paper [13] on the moment-based SDP approach to stochastic chemical kinetics. Many of the ideas it contains can also be found in the other original papers [56, 39, 24]. However, the present chapter also contains several novel elements:

- The systematic use of reaction invariants to further constrain the SDPs and thus obtain tighter bounds for systems for which invariants exist (Section 3.2.3)
- The observation that the most basic instance of the moment-based SDP implicitly relaxes the integrality requirement, and that it is possible to selectively enforce this requirement to obtain tighter bounds (Sections 3.3.4, 3.4.2, 3.6.2, and 3.8.2)
- A method of calculating the probability that the steady-state distribution assigns to a particular interval, which requires solving just one SDP and does not resort to enumeration of the reachable states (Section 3.7)
- The observation that the above method can be generalized to calculate bounds on the probability that the steady-state distribution assigns to an arbitrary basic semi-algebraic set (Section 3.9)

Furthermore, we give additional examples to demonstrate the capabilities and weaknesses of the method. In particular, we include an example reaction system which is considerably more complex than those appearing in the previous papers [13, 56, 39, 24] in that it has 10 different species, 14 reactions, and 74 billion reachable states.

3.2 Mathematical Background

3.2.1 Mathematical Notation

Throughout this thesis, the symbol \mathbb{N} will be used to denote the set of natural numbers $\{0, 1, 2, ...\}$, the symbol \mathbb{Z} will be used to denote the integers $\{\ldots, -2, -1, 0, 1, 2, \ldots\}$, and \mathbb{R} will be used to denote the real numbers. Bold symbols will be used to represent vectors and matrices. The dimensions of these vectors and matrices will be specified as they are introduced. The vector $\mathbf{e}_i = (0, \ldots, 1, \ldots, 0)$ is the *i*th coordinate vector, in which all components are zero, except the *i*th component, which is 1. Angular brackets " $\langle \cdot \rangle$ " will be used to denote an "expected value" or mean of a random variable. The meanings of all other symbols should be clear from the context.

3.2.2 Stochastic Chemical Kinetics Notation

Consider a stochastic chemical kinetic system with N distinct chemical species and R reactions. The state of the system at time t is described by the random vector $\mathbf{X}(t) = (X_1(t), \ldots, X_N(t)) \in \mathbb{N}^N$, where $X_i(t) \ge 0$ is the count of molecules of species *i* present.

The state changes with the occurrence of each reaction. For example, if $\mathbf{s}_r \in \mathbb{Z}^N$ is the vector of stoichiometric coefficients of reaction r, and the system is in state $\mathbf{x} \in \mathbb{N}^N$, then an occurrence of reaction r takes the system to state $\mathbf{x} + \mathbf{s}_r$. By chaining together multiple reactions, a system initially in state $\mathbf{X}(0) \equiv \mathbf{x}_0 \in \mathbb{N}^N$ can reach many possible states – sometimes infinitely many. Let this set of *reachable states* be denoted with the symbol \mathcal{X} . A generic member of this set will be denoted $\mathbf{x} \in \mathcal{X}$.

3.2.3 Invariants and Independent Species

The stoichiometry matrix for the system is constructed by bringing together the stoichiometry vectors: $\mathbf{S} \equiv [\mathbf{s}_1 \dots \mathbf{s}_R] \in \mathbb{Z}^{N \times R}$. Often, this matrix will have a non-trivial left null space. Let $\{\mathbf{b}_1, \dots, \mathbf{b}_L\} \subset \mathbb{R}^N$ be a basis for this left null space. It

can be shown that each of these vectors corresponds to an *invariant* of the reaction system[20] - i.e., some linear combination of molecular counts that is constant with time. In particular,

$$\mathbf{b}_{j}^{\mathrm{T}}\mathbf{X}(t) = \mathbf{b}_{j}^{\mathrm{T}}\mathbf{x}_{0}, \quad \forall j \in \{1, \dots, L\}, \quad \forall t \ge 0.$$
(3.1)

The previous papers [56, 39, 24] on moment-based SDPs in stochastic chemical kinetics do not give a systematic method for dealing with reaction invariants. Most of the examples appearing in these papers do not have reaction invariants. Those that do are handled on an ad hoc basis (see Example 2 in Ghusinga et al[24]). Such ad hoc solutions may not always be so obvious. Furthermore, if you were to take a system with invariants and then apply the basic SDP paradigm described in those papers without some modification to account for the invariants, the resulting bounds would likely be poor. The reason for this, roughly speaking, is that the SDP would be optimizing over distributions that are not only supported on the reachable states $\mathbf{x} \in \mathcal{X} \subset \mathbb{N}^N$, but also states that are not reachable, that is $\mathbf{x} \in \mathbb{N}^N$ such that $\mathbf{x} \notin \mathcal{X}$. So having a systematic method for dealing with invariants is important. We will do this by defining a set of *independent species*.

If we set

$$\mathbf{B} \equiv \begin{bmatrix} \mathbf{b}_{1}^{\mathrm{T}} \\ \vdots \\ \mathbf{b}_{L}^{\mathrm{T}} \end{bmatrix} \in \mathbb{R}^{L \times N}, \tag{3.2}$$

then Equation (3.1) can be expressed concisely as

$$\mathbf{BX}(t) = \mathbf{Bx}_0, \quad \forall t \ge 0. \tag{3.3}$$

These equations imply that not all molecular counts X_1, \ldots, X_N can vary independently. To see this, let $\tilde{\mathbf{B}} \in \mathbb{R}^{L \times L}$ be the matrix obtained by concatenating L linearly independent columns of \mathbf{B} , and let $\tilde{\mathbf{X}}(t) \in \mathbb{N}^L$ be the vector of the corresponding components of $\mathbf{X}(t)$. Similarly, let $\hat{\mathbf{B}} \in \mathbb{R}^{L \times \hat{N}}$ be the matrix obtained by concatenating the remaining $N - L \equiv \hat{N}$ columns of \mathbf{B} , and let $\hat{\mathbf{X}}(t) \in \mathbb{N}^{\hat{N}}$ be the vector of the vector vector of the vector vecto

corresponding components of $\mathbf{X}(t)$. Then, Equation (3.3) can be rewritten as

$$\tilde{\mathbf{B}}\tilde{\mathbf{X}}(t) + \hat{\mathbf{B}}\hat{\mathbf{X}}(t) = \tilde{\mathbf{B}}\tilde{\mathbf{x}}_0 + \hat{\mathbf{B}}\hat{\mathbf{x}}_0, \quad \forall t \ge 0.$$
(3.4)

By construction, $\tilde{\mathbf{B}}$ is invertible, so if $\hat{\mathbf{X}}(t)$ is known, this equation can be solved for $\tilde{\mathbf{X}}(t)$:

$$\tilde{\mathbf{X}}(t) = \tilde{\mathbf{x}}_0 + \tilde{\mathbf{B}}^{-1} \hat{\mathbf{B}}(\hat{\mathbf{x}}_0 - \hat{\mathbf{X}}(t)), \quad \forall t \ge 0.$$
(3.5)

Thus, specifying $\hat{\mathbf{X}}$ is enough to specify the state of the system. We can think of the chemical species whose molecular counts are specified in the vector $\hat{\mathbf{X}}$ as being the independent species, and those whose counts are specified in $\tilde{\mathbf{X}}$ as being the dependent species. In general, there will be several possible ways to pick L linearly independent columns of \mathbf{B} . This means that we have some flexibility in choosing which species to treat as independent.

Every full-dimensional reachable state $\mathbf{x} \in \mathcal{X} \subset \mathbb{N}^N$ has a corresponding *reduced* reachable state, $\hat{\mathbf{x}} \in \mathbb{N}^{\hat{N}}$, obtained by selecting the counts of the independent species from \mathbf{x} . We will denote the set of all these reduced reachable states as $\hat{\mathcal{X}} \subset \mathbb{N}^{\hat{N}}$. Similarly, for every stoichiometry vector $\mathbf{s}_r \in \mathbb{Z}^N$, there is a corresponding reduced stoichiometry vector $\hat{\mathbf{s}}_r \in \mathbb{Z}^{\hat{N}}$, obtained by selecting the components of \mathbf{s}_r corresponding to the independent species.

Working in the reduced state space is computationally convenient because it focuses attention on the variables in the stochastic chemical kinetic system that are actually independent and can thus reduce the dimension of the problems we want to solve. For the sake of brevity, in what follows, we will often loosely refer to the reduced state as simply the "state". That we are in fact referring to the reduced state should be clear from the context.

3.2.4 The Chemical Master Equation

Because of the stochastic nature of the system, there is some uncertainty as to the (reduced) state at time t, and we express this uncertainty by assigning a probability
$\Pr(\hat{\mathbf{X}}(t) = \hat{\mathbf{x}}, t) \equiv P(\hat{\mathbf{x}}, t)$ to each of the reachable states $\hat{\mathbf{x}} \in \hat{\mathcal{X}}$. This probability distribution $P(\cdot, t)$ changes over time according to the chemical master equation (CME):

$$\frac{dP}{dt}(\hat{\mathbf{x}},t) = \sum_{r=1}^{R} [P(\hat{\mathbf{x}} - \hat{\mathbf{s}}_r, t) a_r(\hat{\mathbf{x}} - \hat{\mathbf{s}}_r) - P(\hat{\mathbf{x}}, t) a_r(\hat{\mathbf{x}})],$$

$$\forall \hat{\mathbf{x}} \in \hat{\mathcal{X}},$$
(3.6)

where a_r is the "propensity function" of reaction r. The details of this propensity function are described in Higham[31]. However, we want to point out two things: first, $a_r(\cdot)$ is always a polynomial in $\hat{\mathbf{x}}$; second, a_r is proportional to a rate constant c_r for reaction r. This c_r is not necessarily the same as the macroscopic rate constant k_r one would use in deterministic chemical kinetics, but there is a connection between the two constants. See Higham[31] and Gillespie[26] for details.

Note that the CME holds for *all* reachable states $\hat{\mathbf{x}} \in \hat{\mathcal{X}}$. So it is not just a single equation but a whole system of equations. This system can be written concisely as

$$\frac{d\mathbf{p}}{dt}(t) = \mathbf{G}\mathbf{p}(t),\tag{3.7}$$

where **G** is a time-invariant (infinitesimal generator) matrix whose elements are linked to the propensity functions, and **p** is a vector of probabilities with one component for each $\hat{\mathbf{x}} \in \hat{\mathcal{X}}$. While this equation is conceptually simple, there is a problem: as we suggested earlier, there is often a huge number of reachable states $\hat{\mathbf{x}} \in \hat{\mathcal{X}}$. This means that the vector **p** can have a very large (or even infinite) dimension, with **G** being correspondingly large. The result is that it is impractical to solve Equation (3.7) directly for stochastic chemical kinetic systems of any appreciable size.

3.2.5 Moments in Stochastic Chemical Kinetics

The probability distribution $P(\cdot, t)$ can be characterized by its moments. In particular, for any multi-index $\mathbf{j} = (j_1, \ldots, j_{\hat{N}}) \in \mathbb{N}^{\hat{N}}$ we have a moment $\mu_{\mathbf{j}}(t)$ defined as

$$\mu_{\mathbf{j}}(t) \equiv \sum_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}} \hat{\mathbf{x}}^{\mathbf{j}} P(\hat{\mathbf{x}}, t), \qquad (3.8)$$

where the sum is over the set $\hat{\mathcal{X}}$ of all reachable states, and $\hat{\mathbf{x}}^{\mathbf{j}} = \prod_{k=1}^{\hat{N}} \hat{x}_k^{j_k}$ is a monomial. The *order* of the moment $\mu_{\mathbf{j}}$ is defined as the sum $|\mathbf{j}| \equiv \sum_{k=1}^{\hat{N}} j_k$. Notice that the zeroth-order moment $\mu_{\mathbf{0}}(t)$ indexed by $\mathbf{0} = (0, \dots, 0)$ is simply the sum of probabilities across all reachable states, so that $\mu_{\mathbf{0}}(t) = 1$ for all times t.

A nice feature of moments is that, using just the low-order moments, we can express several quantities of interest that effectively summarize the distribution $P(\cdot, t)$. For example, the first-order moment $\mu_{\mathbf{e}_i}(t)$ indexed by $\mathbf{e}_i = (0, \ldots, 1, \ldots, 0)$ is the mean molecular count for independent species $i \in \{1, \ldots, \hat{N}\}$ at time t:

$$\mu_{\mathbf{e}_{i}}(t) \equiv \sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}} \hat{\mathbf{x}}^{\mathbf{e}_{i}} P(\hat{\mathbf{x}}, t) = \sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}} \hat{x}_{i} P(\hat{\mathbf{x}}, t) = \langle \hat{X}_{i}(t) \rangle.$$
(3.9)

The first-order moments can also be used with Equation (3.5) to express the mean molecular count for each dependent species $k \in \{1, \ldots, L\}$. In particular, if we let $\beta_{k,j}$ denote the element in the *k*th row and *j*th column of the matrix $\tilde{\mathbf{B}}^{-1}\hat{\mathbf{B}}$, then we have

$$\langle \tilde{X}_{k}(t) \rangle = \sum_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}} \left(\tilde{x}_{0,k} + \sum_{j=1}^{N} \beta_{k,j} \hat{x}_{0,j} - \sum_{j=1}^{N} \beta_{k,j} \hat{x}_{j} \right) P(\hat{\mathbf{x}}, t),$$

$$= \left(\tilde{x}_{0,k} + \sum_{j=1}^{\hat{N}} \beta_{k,j} \hat{x}_{0,j} \right) - \sum_{j=1}^{\hat{N}} \beta_{k,j} \mu_{\mathbf{e}_{j}}(t).$$

$$(3.10)$$

Coming to the second-order moments, we see that $\mu_{2\mathbf{e}_i}(t)$ is equal to $\langle \hat{X}_i^2(t) \rangle$. So, $\mu_{\mathbf{e}_i}(t)$ and $\mu_{2\mathbf{e}_i}(t)$ can be used together to compute the variance in the count of molecules of independent species *i* at time *t*:

$$\sigma_i^2(t) \equiv \langle \hat{X}_i^2(t) \rangle - \langle \hat{X}_i(t) \rangle^2 = \mu_{2\mathbf{e}_i}(t) - \mu_{\mathbf{e}_i}^2(t).$$
(3.11)

Similarly, the moments can be used to compute covariances between independent species i and j:

$$\sigma_{i,j}^{2}(t) \equiv \langle \hat{X}_{i}(t)\hat{X}_{j}(t)\rangle - \langle \hat{X}_{i}(t)\rangle\langle \hat{X}_{j}(t)\rangle$$

= $\mu_{\mathbf{e}_{i}+\mathbf{e}_{j}}(t) - \mu_{\mathbf{e}_{i}}(t)\mu_{\mathbf{e}_{j}}(t).$ (3.12)

The appeal of working with moments is that they allow us to bypass the problem of high dimensionality that we encountered in Equation (3.7). We give up a complete

description of the probability distribution $P(\cdot, t)$ in the terms of the high-dimensional vector $\mathbf{p}(t)$ in favor of a summary description in terms of its low-order moments. In principle, this trade-off allows us to compute properties of stochastic chemical kinetic systems for which solving the CME more directly is computationally intractable.

3.2.6 The Closure Problem

As described by Smadbeck and Kaznessis[64], Sotiropoulos and Kaznessis[65], and C.S. Gillespie[25], the CME can be used to derive a system of linear ordinary differential equations describing how the moments of the distribution $P(\cdot, t)$ change over time. For reaction systems containing at most first-order (i.e., unimolecular) reactions, things work out nicely: we can pick an arbitrary $m \in \mathbb{N}$, and construct the ODE describing how the moments up to order m change over time:

$$\frac{d\boldsymbol{\mu}_L}{dt}(t) = \mathbf{A}_L \boldsymbol{\mu}_L(t), \qquad (3.13)$$

where $\boldsymbol{\mu}_L(t)$ is a vector of "low-order" moments order up to order m, and \mathbf{A}_L is a constant matrix. However, if the reaction system contains any reactions of order q > 1 (e.g., bimolecular reactions), then the ODE becomes

$$\frac{d\boldsymbol{\mu}_L}{dt}(t) = \mathbf{A}_L \boldsymbol{\mu}_L(t) + \mathbf{A}_H \boldsymbol{\mu}_H(t), \qquad (3.14)$$

where $\mu_H(t)$ is a vector of "high-order" moments, order m+1 to $m+q-1 \equiv M$. So the time derivatives of the low-order moments depend on high-order moments. This is the infamous "closure problem". It is unclear how to solve such a dynamic system.

The closure problem also frustrates even a relatively simple steady-state analysis. What we'd like to do is set the left-hand side of Equation (3.14) equal to zero

$$\mathbf{0} = \mathbf{A}_L \boldsymbol{\mu}_{L,\mathrm{ss}} + \mathbf{A}_H \boldsymbol{\mu}_{H,\mathrm{ss}},\tag{3.15}$$

and solve for the steady-state moments $\boldsymbol{\mu}_{L,ss}$ and $\boldsymbol{\mu}_{H,ss}$ of the steady-state probability distribution $P_{ss}(\cdot) \equiv \lim_{t \to +\infty} P(\cdot, t)$. Assuming we could calculate the vector $\boldsymbol{\mu}_{L,ss}$, we could extract the steady-state values of $\langle \hat{X}_i \rangle_{ss} \equiv \mu_{\mathbf{e}_i,ss}$ and $(\sigma_i^2)_{ss} \equiv \mu_{2\mathbf{e}_i,ss} - \mu_{\mathbf{e}_i,ss}^2$ for each independent species *i*. The trouble is that Equation (3.15) is under-determined: it has more unknowns than linearly independent equations. Even if we leverage our a priori knowledge of probability distributions and set $\mu_{\mathbf{0},ss} = 1$, one can show there are still more unknowns than linearly independent equations. This means that the system has infinitely many solutions, and we can't simply solve for the steady-state moments $\mu_{L,ss}$ and $\mu_{H,ss}$.

3.2.7 Closure Scheme Approximations

As described in the introduction, various authors have proposed approximations to deal with the closure problem. For example, C. S. Gillespie[25] assumes that the unknown underlying distribution is normal; Smadbeck and Kaznessis[64] assume that it has maximum entropy, given fixed values for the low-order moments. While these schemes differ in their details, they are all fundamentally the same in that they assume some functional relationship between the low- and high-order moments:

$$\boldsymbol{\mu}_H \equiv \mathbf{f}(\boldsymbol{\mu}_L) \tag{3.16}$$

This function need not be explicit. For example, the maximum entropy approach requires solving an optimization problem to evaluate \mathbf{f} . Applying the assumed relationship to (3.14) gives the closed system

$$\frac{d\boldsymbol{\mu}_L}{dt}(t) = \mathbf{A}_L \boldsymbol{\mu}_L(t) + \mathbf{A}_H \mathbf{f}(\boldsymbol{\mu}_L(t)), \qquad (3.17)$$

which can be integrated numerically.

Similarly, applying the assumed relationship to (3.15) gives

$$\mathbf{0} = \mathbf{A}_L \boldsymbol{\mu}_{L,\mathrm{ss}} + \mathbf{A}_H \mathbf{f}(\boldsymbol{\mu}_{L,\mathrm{ss}}). \tag{3.18}$$

The idea is that if the function **f** is chosen appropriately, Equation (3.18), together with $\mu_{0,ss} = 1$, will have a unique solution.

While there are arguments justifying the various closure schemes, ultimately, they are all approximations, with no bounds on the errors they introduce. This is dissatisfying from a theoretical perspective. Furthermore, from a practical perspective, it means that one must check the validity of any results obtained through closure scheme approximations by also applying a second, independent method to the problem, such as Gillespie's Stochastic Simulation Algorithm (SSA) [26]. This is less than ideal. If we have to check the closure scheme results against SSA, why not just run SSA in the first place, and be done with it? We'd prefer a method which retains the computational benefits of moment-based methods, but which also stands on its own and guarantees its own accuracy.

3.3 The Bounding Method

In this section, we describe a fundamentally different approach to dealing with the closure problem. Instead of making an approximation to compute a single number for each of the quantities of interest (means, variances, and so on), we will compute mathematically rigorous bounds on their values.

Admittedly, knowing bounds on the quantities of interest is less useful than knowing the quantities themselves. However, guaranteed bounds on these quantities are arguably more useful than estimates of their values with no bounds on the error. Also, as we will see later, it is sometimes possible to calculate bounds so tight that they effectively give us the desired quantities.

In this chapter, we will consider only the relatively simple steady-state analysis. The extension to the dynamic analysis is discussed in Chapter 4.

3.3.1 The Paradigm

Suppose that we have a generic stochastic chemical kinetic system, characterized by an initial state vector $\mathbf{x}_0 \in \mathbb{N}^N$, a stoichiometry matrix $\mathbf{S} \in \mathbb{Z}^{N \times R}$, and a vector of rate constants $\mathbf{c} \in \mathbb{R}^R$. Assume that there is at least one reaction with order greater than one, so that this system exhibits the closure problem when subjected to a moment analysis. Suppose that we have analyzed **S** to construct an invariant matrix $\mathbf{B} \in \mathbb{R}^{L \times N}$, as described in Section 3.2.3. Suppose further that have identified the $\hat{N} = N - L$ chemical species we wish to treat as independent and constructed the matrices $\hat{\mathbf{B}} \in \mathbb{R}^{L \times \hat{N}}$ and $\tilde{\mathbf{B}} \in \mathbb{R}^{L \times L}$. Finally, suppose that we have chosen a value of $m \in \mathbb{N}$ and constructed the matrices \mathbf{A}_L and \mathbf{A}_H described in Section 3.2.6. We are interested in analyzing the properties of the steady-state distribution(s) of this stochastic chemical kinetic system.

Consider the problem of bounding $\langle \hat{X}_i \rangle_{ss}$, the mean count of molecules of independent species *i* at steady state. What we'd like to do is calculate two numbers $\langle \hat{X}_i \rangle_{ss}^L$ and $\langle \hat{X}_i \rangle_{ss}^U$ such that

$$\langle \hat{X}_i \rangle_{\rm ss}^L \le \langle \hat{X}_i \rangle_{\rm ss} \le \langle \hat{X}_i \rangle_{\rm ss}^U$$
 (3.19)

is guaranteed.

How could we calculate the upper bound, $\langle \hat{X}_i \rangle_{ss}^U$? Suppose that we write down several mathematical constraints that the moment vector $\boldsymbol{\mu}_{ss}$ of any steady-state distribution supported on $\hat{\mathcal{X}}$ must necessarily satisfy – for example, $\mathbf{0} = \mathbf{A}_L \boldsymbol{\mu}_{L,ss} + \mathbf{A}_H \boldsymbol{\mu}_{H,ss}$ from Section 3.2.6. Suppose that we then optimize over the set of vectors $\boldsymbol{\mu}$ satisfying these necessary conditions, searching for the vector $\boldsymbol{\mu}$ with the greatest component $\boldsymbol{\mu}_{\mathbf{e}_i}$:

$$\mu_{\mathbf{e}_{i}}^{*} \equiv \max_{\boldsymbol{\mu}} \quad \mu_{\mathbf{e}_{i}}$$
s.t. $\boldsymbol{\mu}$ satisfies necessary steady-state (3.20)
moment conditions.

Then, because every vector $\boldsymbol{\mu}_{ss}$ corresponding to a steady-state distribution supported on $\hat{\mathcal{X}}$ satisfies the necessary steady-state moment conditions, $\boldsymbol{\mu}_{ss}$ is a feasible point for this optimization problem. It follows that the optimal value $\mu_{\mathbf{e}_i}^*$ is an upper bound on $\mu_{\mathbf{e}_i,ss} = \langle \hat{X}_i \rangle_{ss}$, and we can set $\mu_{\mathbf{e}_i}^* \equiv \langle \hat{X}_i \rangle_{ss}^U$.

Following the same reasoning, we can calculate a lower bound on $\langle \hat{X}_i \rangle_{ss}$ by *minimizing* over the set of vectors $\boldsymbol{\mu}$ satisfying the necessary steady-state moment conditions. It remains to be seen whether or not constructing and solving these optimization problems is computationally tractable. The sole point of this section is a theoretical one: *if* we could solve these optimization problems, we would have our bounds.

3.3.2 Necessary Steady-State Moment Conditions

What exactly are the necessary steady-state moment conditions appearing in Problem (3.20)? We've already identified one:

$$\mathbf{0} = \mathbf{A}_L \boldsymbol{\mu}_L + \mathbf{A}_H \boldsymbol{\mu}_H. \tag{3.21}$$

If we let $\mathbf{A} \equiv [\mathbf{A}_L \ \mathbf{A}_H]$ and $\boldsymbol{\mu}_M = (\boldsymbol{\mu}_L, \boldsymbol{\mu}_H)$, this can be written concisely as

$$\mathbf{0} = \mathbf{A}\boldsymbol{\mu}_M,\tag{3.22}$$

where the subscript M indicates that the vector μ_M contains all moments up to order M (see Section 3.2.6).

The second, relatively obvious necessary condition is that the zeroth moment must be one:

$$\mu_0 = 1. \tag{3.23}$$

Next, we might reason that since all molecular counts must be nonnegative, all mean molecular counts must also be nonnegative. With Equations (3.9) and (3.10), this leads to

$$\mu_{\mathbf{e}_j} \ge 0, \quad \forall j \in \{1, \dots, \hat{N}\} \tag{3.24}$$

for the independent species and

$$\left(\tilde{x}_{0,k} + \sum_{j=1}^{\hat{N}} \beta_{k,j} \hat{x}_{0,j}\right) - \sum_{j=1}^{\hat{N}} \beta_{k,j} \mu_{\mathbf{e}_j} \ge 0, \quad \forall k \in \{1, \dots, L\}$$
(3.25)

for the dependent species.

Also, recalling our expressions for the variances and covariances given by Equa-

tions (3.11) and (3.12), respectively, we can construct a covariance matrix, which we know from probability theory must be positive semidefinite:

$$\begin{bmatrix} \mu_{\mathbf{e}_{1}+\mathbf{e}_{1}}-\mu_{\mathbf{e}_{1}}\mu_{\mathbf{e}_{1}} & \dots & \mu_{\mathbf{e}_{1}+\mathbf{e}_{\hat{N}}}-\mu_{\mathbf{e}_{1}}\mu_{\mathbf{e}_{\hat{N}}} \\ \vdots & \ddots & \vdots \\ \mu_{\mathbf{e}_{\hat{N}}+\mathbf{e}_{1}}-\mu_{\mathbf{e}_{\hat{N}}}\mu_{\mathbf{e}_{1}} & \dots & \mu_{\mathbf{e}_{\hat{N}}+\mathbf{e}_{\hat{N}}}-\mu_{\mathbf{e}_{\hat{N}}}\mu_{\mathbf{e}_{\hat{N}}} \end{bmatrix} \succeq \mathbf{0}.$$
(3.26)

Condition (3.26) is the first of several matrix inequalities that will appear in this chapter. It must be understood that the generic matrix inequality $\mathbf{M} \succeq \mathbf{0}$ does not imply the nonnegativity of each element of the matrix $\mathbf{M} \in \mathbb{R}^{n \times n}$. Rather, it states that all of the eigenvalues of \mathbf{M} are nonnegative. This is equivalent to the statement that $\mathbf{q}^{\mathrm{T}}\mathbf{M}\mathbf{q} \ge 0$ for all $\mathbf{q} \in \mathbb{R}^{n}$. In this thesis, we will only state matrix inequalities for symmetric real-valued matrices.

It turns out that Necessary Conditions (3.24) - (3.26) are all special instances of a more general set of conditions that the moments of any steady-state distribution must satisfy. These more general necessary conditions can be expressed concisely in the form of the following matrix inequalities:

$$\mathbf{M}_{n}^{\mathbf{0}}(\boldsymbol{\mu}) \succeq \mathbf{0}, \tag{3.27}$$

$$\mathbf{M}_{n-1}^{\mathbf{e}_j}(\boldsymbol{\mu}) \succeq \mathbf{0}, \quad \forall j \in \{1, \dots, \hat{N}\},$$
(3.28)

$$\left(\tilde{x}_{0,k} + \sum_{j=1}^{\hat{N}} \beta_{k,j} \hat{x}_{0,j}\right) \mathbf{M}_{n-1}^{\mathbf{0}}(\boldsymbol{\mu}) - \sum_{j=1}^{\hat{N}} \beta_{k,j} \mathbf{M}_{n-1}^{\mathbf{e}_j}(\boldsymbol{\mu}) \succeq \mathbf{0},$$

$$\forall k \in \{1, \dots, L\}.$$
(3.29)

The exact definitions of the matrices $\mathbf{M}_{n}^{\mathbf{0}}(\boldsymbol{\mu})$, $\mathbf{M}_{n-1}^{\mathbf{0}}(\boldsymbol{\mu})$, and $\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\boldsymbol{\mu})$ are unimportant for our present purposes, and have been deferred to Appendix B, where these matrix inequalities are derived. What is important is that each of these matrices is a linear function of the vector $\boldsymbol{\mu}$. This means that each of the necessary conditions (3.27) - (3.29) is a special type of matrix inequality known as a *linear matrix inequal*- *ity* (*LMI*). This is significant because the set of points satisfying a collection of linear matrix inequalities is guaranteed to be a *convex* set. This gives us reason to believe that we may be able to optimize efficiently over this set.

As suggested above, Necessary Conditions (3.24) - (3.26) are subsumed by Necessary Conditions (3.27) - (3.29) in the sense that the former are implied by the latter (see Appendix B). This means that if we specify (3.27) - (3.29) as necessary conditions, it is unnecessary to also specify (3.24) - (3.26). So we may take as our list of necessary conditions (3.22), (3.23), and (3.27) - (3.29).

3.3.3 A Semidefinite Program

If we use Necessary Conditions (3.22), (3.23), and (3.27) - (3.29) in Problem (3.20), in place of the abstract statement " μ satisfies necessary steady-state moment conditions", we obtain

$$\langle \hat{X}_i \rangle_{ss}^U = \max_{\boldsymbol{\mu}} \quad \boldsymbol{\mu}_{\mathbf{e}_i}$$
s.t. $\mathbf{A}\boldsymbol{\mu}_M = \mathbf{0},$

$$\boldsymbol{\mu}_{\mathbf{0}} = 1,$$

$$\mathbf{M}_n^{\mathbf{0}}(\boldsymbol{\mu}) \succeq \mathbf{0}, \qquad (3.30)$$

$$\mathbf{M}_{n-1}^{\mathbf{e}_j}(\boldsymbol{\mu}) \succeq \mathbf{0}, \quad \forall j \in \{1, \dots, \hat{N}\},$$

$$\alpha_k \mathbf{M}_{n-1}^{\mathbf{0}}(\boldsymbol{\mu}) - \sum_{j=1}^{\hat{N}} \beta_{k,j} \mathbf{M}_{n-1}^{\mathbf{e}_j}(\boldsymbol{\mu}) \succeq \mathbf{0},$$

$$\forall k \in \{1, \dots, L\}.$$

Note that we have set $\alpha_k \equiv \tilde{x}_{0,k} + \sum_{j=1}^{\hat{N}} \beta_{k,j} \hat{x}_{0,j}$ for the sake of brevity.

So far, we have been deliberately vague about which moments appear in our vector of decision variables $\boldsymbol{\mu}$. In principle, we could treat every moment in the infinite sequence $\{\mu_{\mathbf{j}} : \mathbf{j} \in \mathbb{N}^{\hat{N}}\}$ as a decision variable. However, this would render numerical solution of Problem (3.30) intractable. So we want to truncate this sequence at some point. It is clear that $\boldsymbol{\mu}$ must include at least those moments up to order M, so that we can meaningfully enforce the constraint $\mathbf{A}\boldsymbol{\mu}_M = \mathbf{0}$. As explained in Appendix B, the LMIs contain moments up to order 2n, where $n \equiv \lceil \frac{M}{2} \rceil$. Since $2n \geq M$, it is sufficient for our decision vector $\boldsymbol{\mu}$ to include only those moments up to order 2n. The key point here is that we have only finitely many decision variables.

While the notation of Problem (3.30) may be somewhat daunting, the problem has nice mathematical properties: the objective function is a (very simple) linear function of the decision variables μ , and the only constraints are linear equalities and linear matrix inequalities. This is a *semidefinite program* (SDP) – a type of convex optimization problem which can be solved efficiently both in theory and in practice [71], using made-for-purpose software such as SeDuMi [66] and MOSEK [2]. So solving the abstract Problem (3.20) is, in fact, computationally tractable. We can readily calculate the desired upper bound on the mean molecular count of independent species *i*. Furthermore, by simply changing the "max" to a "min", we can also calculate $\langle \hat{X}_i \rangle_{ss}^L$. Finally, by replacing the objective function with the final line of Equation (3.10), we could also calculate bounds on the mean molecular counts for each dependent species.

3.3.4 Conservatism in the Bounds

The bounds produced by SDP (3.30) (and its variations) may be *conservative* in the sense that there is a gap between the quantity of interest $\langle \hat{X}_i \rangle_{ss}$ and the bound $\langle \hat{X}_i \rangle_{ss}^U$. There are several sources of this conservatism, but they can all be thought of generally as an omission of necessary conditions from SDP (3.30). This idea will become more clear in the following discussion of the two most important sources of conservatism.

Choosing $m \in \mathbb{N}$

The first necessary condition appearing in SDP (3.30) is the equation $\mathbf{A}\boldsymbol{\mu}_M = \mathbf{0}$. Recall from Section 3.2.6 that this equation sets to zero the time derivatives of all moments up to order $m \in \mathbb{N}$. However, we know that for any true steady-state distribution *all* moments are unchanging with time. Thus, $\mathbf{A}\boldsymbol{\mu}_M = \mathbf{0}$ enforces only a finite subset of an infinite sequence of necessary conditions on the steady-state moments. By omitting the conditions specifying that the moments with order greater than m must have zero time derivatives, we are enlarging the feasible set of SDP (3.30), admitting more moment vectors $\boldsymbol{\mu}$ to the competition. This has the potential to increase the optimal value $\langle \hat{X}_i \rangle_{ss}^U$ of the SDP, leading to conservatism in the bound.

Now, recall that our choice of $m \in \mathbb{N}$ was arbitrary. Given the foregoing argument, it makes sense to choose m as large as possible, reasoning that doing so might lead to tighter bounds. However, there is a trade-off: a large value of m implies that SDP (3.30) becomes large in terms of both the number of constraints and variables. In general, larger problems take longer to solve and are more prone to numerical issues. So there's a balance to be struck in choosing m. Our general strategy in what follows will be to start by choosing a low value of m and then increase it as necessary to improve the quality of the bounds.

Relaxing the Integrality Requirement

Our aim in the Section 3.3.2 was to list necessary conditions for the moments μ of any steady-state distribution supported on the discrete set of reachable states $\hat{\mathcal{X}} \subset \mathbb{N}^{\hat{N}}$. However, what we've actually done is list necessary conditions for the moments μ of any steady-state distribution supported on the polyhedral set

$$\bar{\mathcal{X}} \equiv \left\{ \hat{\mathbf{x}} \in \mathbb{R}^{\hat{N}} : \begin{array}{c} \hat{\mathbf{x}} \ge \mathbf{0}, \\ \tilde{\mathbf{x}}_0 + \tilde{\mathbf{B}}^{-1} \hat{\mathbf{B}} (\hat{\mathbf{x}}_0 - \hat{\mathbf{x}}) \ge \mathbf{0} \end{array} \right\}.$$
(3.31)

In particular, the inequalities defining $\bar{\mathcal{X}}$ are reflected in LMIs (3.28) and (3.29).

By the construction of $\bar{\mathcal{X}}$, we are guaranteed that $\hat{\mathcal{X}}$ is a subset of $\bar{\mathcal{X}}$. This means that Conditions (3.28) and (3.29) are also necessary for the moments μ of any steadystate distribution supported on $\hat{\mathcal{X}}$. That's good – it means that solving SDP (3.30) will, in fact, produce valid bounds, as we intended. The downside to having our necessary conditions based on the superset $\bar{\mathcal{X}}$ is that we are including in the feasible set of SDP (3.30) moment vectors μ corresponding to *nonphysical* distributions – namely, distributions which assign a nonzero probability to states $\hat{\mathbf{x}} \in \bar{\mathcal{X}}$ which imply non-integer molecular counts. This has the potential to introduce undesirable conservatism into the bounds produced by solving SDP (3.30).

Given this potential for undesirable conservatism in the bounds, one might consider adding further necessary conditions to exclude the nonphysical distributions and focus attention on the moments of distributions supported on the discrete set $\hat{\mathcal{X}}$. In principle, this can be done (See Lasserre^[42], Chapter 5). However, recall that the fundamental problem with the CME is that there is usually a very large number of states $\hat{\mathbf{x}} \in \hat{\mathcal{X}}$. If we attempt to add constraints to SDP (3.30) to restrict the support of the unknown steady-state distribution to $\hat{\mathcal{X}}$, we find that we quickly have an intractably large number of constraints. So the problem of high dimensionality that we saw in Section 3.2.4 just manifests itself in a different way. As discussed previously, the whole point of the moment-based analysis of stochastic chemical kinetics is to avoid this problem of high dimensionality. So we will not try to restrict the support of the unknown steady-state distribution to the discrete set $\hat{\mathcal{X}}$. We will content ourselves with Necessary Conditions (3.28) and (3.29), thereby considering the moments of all steady-state distributions supported on the larger set \mathcal{X} . This constitutes a *relaxation* of the integrality requirement. We accept the conservatism that is introduced through this relaxation as the price of computational tractability.

The SDPs described by Ghusinga et al[24], Kuntz et al[39], and Sakurai et al[56] also implicitly relax the integrality requirement. However, there is no discussion of this fact and its implications in those papers.

3.3.5 Scaling

One difficulty with using moments as decision variables in SDP (3.30) is that the moments of a distribution often differ by orders of magnitude. This can make momentbased SDPs ill-conditioned, leading to numerical difficulties[39, 18]. These difficulties can be partially alleviated by appropriately *scaling* the decision variables. Since this scaling procedure is a matter of numerics and is not central to the theory of this chapter, we have deferred it to Appendix B. However, going forward, whenever we speak of solving SDP (3.30), it should be understood that we mean the appropriately scaled version of SDP (3.30). These two SDPs are mathematically equivalent; it's just that the scaled version improves the performance of the solvers.

3.4 Mean-Bounding Examples

The previous section established that we can, in fact, calculate theoretically guaranteed bounds on the mean molecular counts for each species. However, it remains to be seen if these bounds are actually useful.

Ideally, the bounds on any given mean molecular count would be equal, in which case they would tell us the mean value exactly. The second-best scenario is that the bounds are tight, in the sense that they have only a small gap between them, bracketing the unknown mean in a relatively small range. What we really don't want is for there to be a large gap between the bounds, in which case there is still a lot of uncertainty surrounding the value of the mean.

In the examples of this section, we will see that the quality of the bounds depends on the characteristics of the stochastic chemical kinetic system.

3.4.1 Michaelis-Menten

The first example is a variant of the classical Michaelis-Menten system, taken from Smadbeck and Kaznessis[64].

$$S + E \xrightarrow{c_1} S:E$$

$$S:E \xrightarrow{c_3} P + E \qquad (3.32)$$

$$P \xrightarrow{c_4} S$$

In that paper, Smadbeck and Kaznessis calculate approximate values of $\langle S \rangle_{ss}$ and $\langle E \rangle_{ss}$ assuming initial molecular counts S = 10, E = 10, S:E = 0, and P = 0, fixed rate constants $c_2 = c_3 = c_4 = 1 \text{ s}^{-1}$, and c_1 values ranging from 10^{-3} s^{-1} to 10^{-2} s^{-1} .



Figure 3-1: Bounds on the Michaelis-Menten system at steady state. The points marked with circles and stars each correspond to the solution of an SDP and are theoretically guaranteed bounds. The lines interpolated between these points are not guaranteed bounds. They are included just to lead the eye.

We calculated bounds on $\langle S \rangle_{ss}$ and $\langle E \rangle_{ss}$ for these same conditions by repeatedly solving SDP (3.30) and the corresponding minimization problem, taking m = 2. The results are shown in Figure 3-1. In this case, the bounds are so tight that they effectively give the quantities $\langle S \rangle_{ss}$ and $\langle E \rangle_{ss}$ for each value of c_1 . Increasing the value of m in an effort to obtain tighter bounds was deemed unnecessary.

Figure 3-1 is agreement with Smadbeck and Kaznessis's Figure 3A. However, note the difference in what is being plotted: Smadbeck's Figure 3A provides an estimate of $\langle S \rangle_{ss}$ and $\langle E \rangle_{ss}$ with no bound on the error, while Figure 3-1 provides guaranteed bounds on $\langle S \rangle_{ss}$ and $\langle E \rangle_{ss}$. Thus, Figure 3-1 contains more information.

3.4.2 Reversible Dimerization

The second example is a simple reversible dimerization reaction, also taken from Smadbeck [64].

$$2A \underbrace{\stackrel{C_1}{\longleftarrow}}_{C_2} B \tag{3.33}$$



Figure 3-2: Bounds on the Reversible Dimerization system at steady state

Smadbeck calculates approximate values of $\langle A \rangle_{ss}$ assuming initial molecular counts A = 10 and B = 0, a fixed rate constant $c_1 = 1 \text{ s}^{-1}$ and c_2 values ranging from 10^{-3} s^{-1} to 10^3 s^{-1} .

We calculated bounds on $\langle A \rangle_{ss}$ for these same conditions by repeatedly solving SDP (3.30) and the corresponding minimization problem taking m = 2. The results are shown in the top plot of Figure 3-2, along with an analytical solution provided by McQuarrie[47].

The first thing to notice in this plot is that the bounds do, in fact, enclose the analytical solution, supporting the validity of our theory.

The second thing to notice is that while the bounds are very tight for large values of c_2 , there is a non-negligible gap for small values. In an effort to eliminate this gap, we

calculated the bounds again, taking m = 4. While the bounds did tighten somewhat for intermediate values of c_2 , the gap at the far left of the plot stubbornly persisted, suggesting that it is attributable to something fundamental to the formulation of SDP (3.30), and cannot be dealt with by simply increasing m.

This is indeed the case. The gap in Figure 3-2 is an example of the conservatism introduced by relaxing the integrality requirement, as discussed in Section 3.3.4. It is easiest to see this if we examine the limiting case where $c_2 = 0$. In this limit, there is no reverse reaction. This means that the reaction system can only be at steady state if the forward reaction cannot occur. Since the forward reaction involves two molecules of A, there are exactly two situations when it cannot occur: when there are 0 molecules of A and when there is 1 molecule of A. Given the stoichiometry and initial condition for this system, the latter situation is not possible. We can see this by looking at the Invariant Equation (3.5) for this system, which reduces to $x_{\rm A} + 2x_{\rm B} = 10$. Having 1 molecule of A implies 4.5 molecules of B – a nonphysical noninteger number. This means, if $c_2 = 0$, the only possible steady-state distribution for this system is a Dirac distribution where all the probability is concentrated on the point $x_{\rm A} = 0$. This implies perfectly tight bounds: $\langle {\rm A} \rangle_{\rm ss}^U = \langle {\rm A} \rangle_{\rm ss}^L = 0$. However, because we relaxed the integrality requirement in our construction of SDP (3.30), the state with 1 molecule of A is also allowed. This means that there is a second distribution whose moments satisfy all the constraints of SDP (3.30), including the steady-state constraint Ay = 0, namely the Dirac distribution centered on the point $x_{\rm A} = 1$, with corresponding mean $\langle A \rangle_{\rm ss} = 1$. The feasibility of the moments of this distribution pushes the upper bound on the mean up to $\langle A \rangle_{ss}^U = 1$, creating a gap between the bounds. In this way, relaxing the integrality requirement introduces conservatism in the bounds. The situation is similar for $c_2 > 0$, though not amenable to the sort of analysis we have applied here. The gap vanishes as $c_2 \rightarrow +\infty$, because, in this limit, only the reverse reaction is relevant, and the only situation in which the reverse reaction cannot occur is when there are 0 molecules of B (and 10 molecules of A).

As suggested in Section 3.3.4, we consider the conservatism introduced through

relaxing the integrality requirement to be an acceptable price to pay for computational tractability. However, we recognize that in some cases it may be necessary to reduce this conservatism, even at the expense of solving a slightly larger SDP. We can do this by selectively adding LMIs reflecting the integrality requirement to SDP (3.30). For example, by adding a single LMI to SDP (3.30) and repeating the previous analysis, we can obtain the tighter bounds shown in the bottom plot of Figure 3-2. The details of the construction of this LMI can be found in Appendix B.

This discussion of the gap in the bounds and how to deal with it is related to the discussion of "non-unique stationary distributions" in the paper by Kuntz et al[39]. However, Kuntz describes a gap in the bounds only in the context of solving a Linear Program (LP), where each decision variable corresponds to a particular reachable state $\mathbf{x} \in \mathcal{X}$. Thus, his interpretation of the source of the gap and subsequent solution are quite different from our own. He does not address the possibility that such a gap could be observed in solving a moment-based SDP.

3.4.3 Schlögl

The third example is the Schlögl model, also taken from Smadbeck and Kaznessis[64].

$$2X + A \xrightarrow[c_2]{c_2} 3X$$

$$B \xrightarrow[c_3]{c_4} X$$
(3.34)

Species A and B are assumed to be present in large quantities and buffered, so that their counts remain essentially constant with time. This means that X is the only species of interest. The Schlögl model is a favorite toy problem because its steadystate distribution for X is bimodal for some parameter values. Furthermore, there are infinitely many reachable states, so there is no hope of analyzing the problem by directly solving the CME.

Smadbeck and Kaznessis approximate values of $\langle X \rangle_{ss}$, assuming $c_1 x_A = 0.15 \text{ s}^{-1}$, $c_2 = 1.5 \times 10^{-3} \text{ s}^{-1}$, $c_3 x_B = 20 \text{ s}^{-1}$, and c_4 ranging from 2 s⁻¹ to 5 s⁻¹. We calculated bounds on $\langle X \rangle_{ss}$ for these same conditions by repeatedly solving SDP (3.30) and

the corresponding minimization problem, taking m = 2. The resulting bounds were surprising poor, with a large gap for all but the largest values of c_4 . However, we found that by increasing the value of m, we could gradually improve the quality of the bounds. This gradual improvement is shown in Figure 3-3, where we plot the bounds for m = 2, m = 6, and m = 10. For higher values of m, the improvement in the bounds was negligible. At m = 16, we began to encounter numerical difficulties.

The bounds associated with m = 10 are tight for both high and low values of c_4 . However, there is still a gap from about $c_4 = 3 \text{ s}^{-1}$ to $c_4 = 4 \text{ s}^{-1}$. Interestingly, this is roughly the range of c_4 values for which the steady-state distribution is bimodal, given the assumed values of c_1, c_2 and c_3 (see Appendix B). This might suggest that the relative difficulty in obtaining tight bounds for this system is attributable to its capacity to exhibit a bimodal steady-state distribution.

Figure 3-3 is in agreement with Smadbeck and Kaznessis's Figure 5D. However, again, it is important to notice the difference in what is being plotted: Smadbeck and Kaznessis's Figure 5D provides an estimate of $\langle X \rangle_{ss}$ with no bound on the error, while Figure 3-3 provides guaranteed bounds on $\langle X \rangle_{ss}$.

3.4.4 A Larger Example

So far, our examples have been concerned with small toy models from the literature. The fourth example, shown in Figure 3-4, is one of our own creation, designed to show that our bounding method can also be successfully applied to larger models with features that frustrate other methods.

This system has 14 reactions and 10 chemical species. It includes unimolecular and bimolecular reactions, reversible and irreversible reactions, fast and slow reactions (rate constants given in Appendix B), and a reaction cycle. We assumed initial molecular counts A = 100, F = 100, and zero for all other species. With this initial condition, one can calculate (see Appendix B) that the number of reachable states for this system is 74,816,108,146. So directly solving the CME for this system is intractable. Furthermore, because the presence of both fast and slow reactions, studying this system using the SSA algorithm (or any of its exact variants) is computationally



Figure 3-3: Bounds on the Schlögl system at steady state



Figure 3-4: A larger reaction system

inefficient.

As with the previous examples, we performed a sensitivity study, varying the rate constant for $A \rightarrow B + C$ from $c_1 = 1 \times 10^{-3} \text{ s}^{-1}$ to $c_1 = 1 \times 10^3 \text{ s}^{-1}$, while calculating bounds on $\langle A \rangle_{ss}$, $\langle D \rangle_{ss}$, and $\langle E \rangle_{ss}$. The bounds were calculated by repeatedly solving SDP (3.30), taking m = 2 and m = 4. The results are shown in Figure 3-5.

While there are small gaps in the bounds, on the whole, they are relatively tight, effectively giving us the desired quantities $\langle A \rangle_{ss}$, $\langle D \rangle_{ss}$, and $\langle E \rangle_{ss}$ across the range of c_1 values sampled. Thus, we see that the bounding method is not limited to small toy models, but can also be successfully applied to a larger reaction system with features that prevent effective analysis using other methods.

3.5 Variance Bounds

In the previous section, we have demonstrated that we can bound the mean count of molecules of each species at steady-state. In this section, we show that we can also bound the variances in each of these counts.



Figure 3-5: Bounds on the larger reaction system in Figure 3-4 at steady state.

3.5.1 An Upper-Bounding SDP

A natural starting point in attempting to bound the steady-state variance of species i is to replace the objective function of SDP (3.30) with the expression given for σ_i^2 in Equation (3.11). Doing so gives us the following optimization problem:

$$(\sigma_i^2)_{ss}^U \equiv \max_{\boldsymbol{\mu}} \quad \mu_{2\mathbf{e}_i} - \mu_{\mathbf{e}_i}^2$$

s.t. $\mathbf{A}\boldsymbol{\mu}_M = \mathbf{0},$
 $\mu_{\mathbf{0}} = 1,$
 $\mathbf{M}_n^{\mathbf{0}}(\boldsymbol{\mu}) \succeq \mathbf{0},$ (3.35)
 $\mathbf{M}_{n-1}^{\mathbf{e}_j}(\boldsymbol{\mu}) \succeq \mathbf{0}, \quad \forall j \in \{1, \dots, \hat{N}\},$
 $\alpha_k \mathbf{M}_{n-1}^{\mathbf{0}}(\boldsymbol{\mu}) - \sum_{j=1}^{\hat{N}} \beta_{k,j} \mathbf{M}_{n-1}^{\mathbf{e}_j}(\boldsymbol{\mu}) \succeq \mathbf{0},$
 $\forall k \in \{1, \dots, L\}.$

By the same logic we applied in Section 3.3.1, the optimal value of Problem (3.35) is a theoretically guaranteed upper bound on $(\sigma_i^2)_{ss}$.

However, the question remains as to whether or not we can actually solve Problem (3.35) efficiently. Recall that the reason we can solve the mean-bounding Problem (3.30) efficiently is that it is an SDP, with all the nice mathematical properties this entails. Problem (3.35) is nearly an SDP, but not quite – the objective function is a nonlinear function of the decision variables. At first, it may seem that this nonlinearity would frustrate our attempts to solve Problem (3.35) efficiently. However, Problem (3.35) can be reformulated as an SDP in just two steps.

The first step is to introduce a dummy decision variable s and move the nonlin-

earity to the constraints:

$$(\sigma_{i}^{2})_{ss}^{U} = \max_{\boldsymbol{\mu},s} \quad s$$

s.t. $s \leq \mu_{2\mathbf{e}_{i}} - \mu_{\mathbf{e}_{i}}^{2}$
 $\mathbf{A}\boldsymbol{\mu}_{M} = \mathbf{0},$
 $\mu_{\mathbf{0}} = 1,$
 $\mathbf{M}_{n}^{\mathbf{0}}(\boldsymbol{\mu}) \succeq \mathbf{0},$ (3.36)
 $\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\boldsymbol{\mu}) \succeq \mathbf{0}, \quad \forall j \in \{1, \dots, \hat{N}\},$
 $\alpha_{k}\mathbf{M}_{n-1}^{\mathbf{0}}(\boldsymbol{\mu}) - \sum_{j=1}^{\hat{N}} \beta_{k,j}\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\boldsymbol{\mu}) \succeq \mathbf{0},$
 $\forall k \in \{1, \dots, L\}.$

Problems (3.36) and (3.35) are equivalent in the sense that they have the same optimal value.

The second step is based on the following equivalence due to the Schur Complement Lemma [71]:

$$s \le \mu_{2\mathbf{e}_i} - \mu_{\mathbf{e}_i}^2 \quad \iff \quad \begin{bmatrix} \mu_{2\mathbf{e}_i} - s & \mu_{\mathbf{e}_i} \\ \mu_{\mathbf{e}_i} & 1 \end{bmatrix} \succeq \mathbf{0}$$
(3.37)

The inequality on the right is an LMI. Substituting it into Problem (3.36) in place of

 $s \leq \mu_{2\mathbf{e}_i} - \mu_{\mathbf{e}_i}^2$, we obtain an SDP for calculating an upper bound on $(\sigma_i^2)_{ss}$:

$$(\sigma_i^2)_{ss}^U = \max_{\boldsymbol{\mu}, s} \quad s$$

s.t.
$$\begin{bmatrix} \mu_{2\mathbf{e}_i} - s & \mu_{\mathbf{e}_i} \\ \mu_{\mathbf{e}_i} & 1 \end{bmatrix} \succeq \mathbf{0}$$

$$\mathbf{A}\boldsymbol{\mu}_M = \mathbf{0},$$

$$\mu_{\mathbf{0}} = 1,$$

$$\mathbf{M}_n^{\mathbf{0}}(\boldsymbol{\mu}) \succeq \mathbf{0},$$

$$\mathbf{M}_{n-1}^{\mathbf{e}_j}(\boldsymbol{\mu}) \succeq \mathbf{0}, \quad \forall j \in \{1, \dots, \hat{N}\},$$

$$\alpha_k \mathbf{M}_{n-1}^{\mathbf{0}}(\boldsymbol{\mu}) - \sum_{j=1}^{\hat{N}} \beta_{k,j} \mathbf{M}_{n-1}^{\mathbf{e}_j}(\boldsymbol{\mu}) \succeq \mathbf{0},$$

$$\forall k \in \{1, \dots, L\}.$$

(3.38)

3.5.2 A Lower-Bounding SDP?

In the case of the mean-bounding problem, we were able to convert the SDP for obtaining the upper bound $\langle \hat{X}_i \rangle_{ss}^U$ into an SDP for obtaining the lower bound $\langle \hat{X}_i \rangle_{ss}^L$ simply by changing the "max" to a "min". One might suspect that we could do something similar with SDP (3.38). However, the minimization version of SDP (3.38) is meaningless. The reason for this is that the minimization versions of Problems (3.35) and (3.36) are in general *not* equivalent. Thus, the argument in moving from the minimization version of SDP (3.38) to the minimization version of SDP (3.38) breaks down.

We cannot formulate a single SDP that is equivalent to Problem (3.35), because the objective function of that problem is concave, and minimizing a concave function is a nonconvex problem. Sakurai et al[56] point out that we can obtain a relaxed lower bound on the variance by solving two SDPs successively. However, we will not pursue this strategy here. We always have the trivial lower bound of $(\sigma_i^2)_{ss}^L = 0$. Furthermore, from a practical perspective, an upper bound on the variance (a measure of uncertainty) seems more useful than a lower bound.

3.6 Variance-Bounding Examples

We next apply SDP (3.38) to the example models appearing in Section 4.5, bounding the variance of each species for which we previously bounded the mean. For each example model, we assume the same initial molecular counts and rate constants as before.

3.6.1 Michaelis-Menten

Repeatedly solving SDP (3.38) with m = 2 for the Michaelis-Menten system (3.32), we obtain the bounds shown in the top plot of Figure 3-6. The figure also shows the exact variances obtained by solving the CME for this system. Comparing the bounds with the exact variances, we see that the bounds are valid but that they are a bit conservative for intermediate values of c_1 .

Repeating the analysis with m = 4, we obtain the bounds shown in the bottom plot of Figure 3-6. These bounds follow the CME-generated variances so closely that the curves are indistinguishable.

3.6.2 Reversible Dimerization

Repeatedly solving SDP (3.38) with m = 2 for the Reversible Dimerization system (3.33), we obtain the bounds shown in the top plot of Figure 3-7. In this case, we have an analytical solution for the steady-state variance [47], which we've also plotted in Figure 3-7. This analytical solution makes two things clear: first, the bounds obtained through SDP (3.38) are, in fact, valid over-estimators of the true variances; second, the bounds are conservative for intermediate values of c_2 .

Repeating the analysis with m = 4, we obtain the bounds shown in the bottom plot of Figure 3-7. These bounds are tighter for the intermediate values of c_2 , closely following the analytical curve. However, notice that a gap persists at the low end of the c_2 range. This gap can be traced to the conservatism introduced by relaxing the integrality requirement, just like the gap we saw in Figure 3-2, in bounding the mean for this system.



Figure 3-6: Bounds on variances for the Michaelis-Menten system at steady state



Figure 3-7: Bounds on variances for the Reversible Dimerization system at steady state

3.6.3 Schlögl

Repeatedly solving SDP (3.38) with m = 2 for the Schlögl system (3.34), we obtain the bounds shown in the top plot of Figure 3-8. Repeating the analysis with m = 4, we obtain the tighter bounds shown in the bottom plot of Figure 3-8.

The two bounding curves of Figure 3-8 are consistent with Smadbeck's Figure S7.4(b). However, taking the points of Smadbeck's Figure S7.4(b) generated through SSA as the true variance values, even the relatively tight bounds in the bottom panel seem to be rather conservative. In some sense, this isn't surprising, given that we had to choose m = 10 to obtain reasonably tight bounds on the mean for this system (see Figure 3-3). This suggests that we would have to use similarly large values of m to obtain reasonably tight bounds on the variance. Unfortunately, when we attempted to repeat the analysis with m = 6, we encountered numerical issues. So, it seems that this system's capacity to exhibit a bimodal steady-state distribution makes it difficult to obtain tight bounds for the variance as well as the mean. All this being said, we wish to emphasize that while the bounds shown in Figure 3-3 are conservative, they are still valid, theoretically guaranteed bounds.

3.6.4 A Larger Example

Repeatedly solving SDP (3.38) with m = 2 for the larger reaction system in Figure 3-4, we obtain the bounds shown in Figure 3-9. When we attempted to repeat the analysis with larger values of m, we encountered numerical difficulties.

3.7 Bounds on Probability

In the previous sections, we've seen that we can bound the mean and variance of the molecular count of each species in the steady-state probability distribution. In this section, we show that we can also calculate bounds on the probability that the molecular count of independent species *i* lies in a specific interval, $[x_{\min}, x_{\max}]$.



Figure 3-8: Bounds on variances for the Schlögl system at steady state



Figure 3-9: Bounds on variances for the larger reaction system in Figure 3-4 at steady state.

3.7.1 An Upper Bounding SDP

As with the mean and variance bounds, the key to calculating an upper bound on the probability that $\hat{X}_i \in [x_{\min}, x_{\max}]$ is an SDP:

$$P_{\rm ss}(\hat{X}_{i} \in [x_{\min}, x_{\max}])^{U} \equiv$$

$$\max_{\mathbf{y}, \mathbf{z}} z_{\mathbf{0}}$$
s.t. $\mathbf{A}\mathbf{y}_{M} + \mathbf{A}\mathbf{z}_{M} = \mathbf{0}, \quad y_{\mathbf{0}} + z_{\mathbf{0}} = 1,$

$$\mathbf{M}_{n}^{\mathbf{0}}(\mathbf{y}) \succeq \mathbf{0}, \quad \mathbf{M}_{n}^{\mathbf{0}}(\mathbf{z}) \succeq \mathbf{0}$$

$$\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\mathbf{y}) \succeq \mathbf{0}, \quad \mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\mathbf{z}) \succeq \mathbf{0}, \quad \forall j \in \{1, \dots, \hat{N}\},$$

$$\alpha_{k}\mathbf{M}_{n-1}^{\mathbf{0}}(\mathbf{y}) - \sum_{j=1}^{\hat{N}} \beta_{k,j}\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\mathbf{y}) \succeq \mathbf{0}, \quad \forall k \in \{1, \dots, L\},$$

$$\alpha_{k}\mathbf{M}_{n-1}^{\mathbf{0}}(\mathbf{z}) - \sum_{j=1}^{\hat{N}} \beta_{k,j}\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\mathbf{z}) \succeq \mathbf{0}, \quad \forall k \in \{1, \dots, L\},$$

$$-\mathbf{M}_{n-2}^{2\mathbf{e}_{i}}(\mathbf{z}) + (x_{\max} + x_{\min})\mathbf{M}_{n-2}^{\mathbf{e}_{i}}(\mathbf{z}) - x_{\min}x_{\max}\mathbf{M}_{n-2}^{\mathbf{0}}(\mathbf{z}) \succeq \mathbf{0},$$

$$(3.39)$$

The derivation of SDP (3.39) requires a lengthy explanation and is deferred to Appendix B.

While SDP (3.39) is certainly more complicated than the SDPs we've seen previously, it is just a variation on the same theme: again we have an equation involving the **A** matrix, corresponding to the steady-state condition; again we have the specification that the sum of probabilities must equal one; and again we have LMIs reflecting constraints on the support of the distribution(s). The primary difference is that we now have two moment vectors **y** and **z**, corresponding to two different distributions. The other new element is the presence of the final LMI, which reflects the interval of interest $[x_{\min}, x_{\max}]$.

To be clear, SDP (3.39) can be used to calculate upper bounds on probabilities only for an independent species $i \in \{1, ..., \hat{N}\}$. In principle, we could derive a similar SDP for calculating bounds for the dependent species. However, if you want to calculate a bound for some dependent species $k \in \{1, ..., L\}$, it is probably simplest to just choose your independent species differently, so that the species of interest is in that set. Then, you can use SDP (3.39) without any modification or added complexity.

In the hopes of improving the quality of the bound produced by SDP (3.39), one might consider adding LMIs to restrict the support of the underlying probability distribution to the integer values in the interval $[x_{\min}, x_{\max}]$. However, we found that doing so had only a negligible effect.

3.7.2 A Lower Bounding SDP

One can also derive an SDP for calculating a lower bound on the probability that $\hat{X}_i \in [x_{\min}, x_{\max}]$. However, our experience with this SDP is that it usually gives a trivial bound of $P_{\rm ss}(\hat{X}_i \in [x_{\min}, x_{\max}])^L = 0$ unless the interval $[x_{\min}, x_{\max}]$ is very wide relative to the spread of the true steady-state probability distribution. So it is of limited practical value.

3.7.3 Bounds on a Histogram

The ability to calculate an upper bound on the probability that the molecular count of species *i* lies between x_{\min} and x_{\max} is interesting in itself. But by chaining together several such calculations, we can obtain something even more interesting: upper bounds on a *histogram* describing the steady-state probability distribution. We simply pick a set of intervals $[x_{\min}^{(j)}, x_{\max}^{(j)}]$ and then solve SDP (3.39) for each interval.

The choice of the intervals $[x_{\min}{}^{(j)}, x_{\max}{}^{(j)}]$ is arbitrary. However, there are a few things to keep in mind when making this choice. First and most obviously, we will want to choose the intervals so that they cover the most probable values of \hat{X}_i . Analysis of the reaction stoichiometry as well as our bounds on $\langle \hat{X}_i \rangle_{ss}$ and $(\sigma_i^2)_{ss}$ may be helpful here. Second, the narrower the intervals, the better our bounds will resolve the features of the unknown probability distribution (e.g., multiple modes). Third, the unknown probability distribution is nonzero only on nonnegative integer values of \hat{x}_i . This suggests that we should choose our intervals to be singletons $[x^{(j)}, x^{(j)}]$, where $x_{\min}{}^{(j)} = x_{\max}{}^{(j)} = x{}^{(j)}$ and each $x{}^{(j)}$ is a nonnegative integer molecular count. This singleton interval approach makes sense in the case when the range of molecular counts that we care about is small. However, recall that we have to solve an SDP for each interval. So, if the range of molecular counts that we care about is relatively large, having an interval for each count may be impractical, because it requires us to solve a large number of SDPs. In this case, we will select our intervals so that each covers several molecular counts (i.e., $x_{\min}^{(j)} < x_{\max}^{(j)}$). Certainly, we lose some resolution in doing this, but it may be necessary to maintain computational tractability. In the following examples, we will use both singleton and non-singleton intervals.

The method that we have described here is related to the linear programming approach for calculating bounds on marginal distributions described by Kuntz et al[39]. However, there are fundamental differences between the two methods. Kuntz's approach requires enumeration of part of the set of reachable states, and so there are theoretical concerns about its scalability. In contrast, our method does not require enumeration of any part of the set of reachable states. Furthermore, to compute an upper bound on the probability associated with a particular reachable state, Kuntz's method requires the solution of an SDP followed by an LP, with ad hoc selection of several parameters, including the "weight vector", "state space truncation parameter", and the order of the moment used in the LP. In contrast, our method requires the solution to just one SDP and has no analogous parameters.

3.8 Probability-Bounding Examples

We again revisit each of the example models we've seen previously. This time, instead of varying one of the rate constants, we will fix the rate constants and vary the interval $[x_{\min}, x_{\max}]$, solving SDP (3.39) for each interval. Doing so, we obtain upper bounds on a histogram for any steady-state probability distribution. We will see that while these bounds often substantially over-estimate the true histogram, they still give a reasonable picture of its shape.

3.8.1 Michaelis-Menten

For the Michaelis-Menten Model (3.32), we assume the same values for the rate constants and initial counts stated in Section 3.4.1 – the one exception being c_1 , which we now hold fixed at $c_1 = 1 \text{ s}^{-1}$. We aim to obtain bounds on the molecular count histogram for species S. Given the stoichiometry of the reaction system and the initial condition, we see that the molecular count of S can take on at most 11 values, $x_S = 0, \ldots, 10$. Accordingly, we will pick our intervals as the singletons $[0, 0], [1, 1], \ldots [10, 10]$. Solving SDP (3.39) with m = 2 and m = 4 for each of these "intervals" gives the upper bounding histograms shown in Figure 3-10.

For comparison, each of the plots in Figure 3-10 also shows the histogram of $P(x_S)$ values for the true steady-state probability distribution, obtained through solving the CME. For many values of x_S , the bound $P_{ss}(x_S)^U$ greatly exceeds the true probability $P_{ss}(x_S)$. However, notice that the shape of the upper bounding histogram closely follows that of the true histogram.

3.8.2 Reversible Dimerization

For the Reversible Dimerization Model (3.33), we assume the same values for the rate constants and initial counts stated in Section 3.4.2 – the one exception being c_2 , which we now hold fixed at $c_2 = 1 \text{ s}^{-1}$. We aim to obtain bounds on the molecular count histogram for species A. Given the stoichiometry of the reaction system and the initial condition, we see that the molecular count of A is bounded above by 10 and bounded below by 0. So again we will pick our intervals as the singletons $[0,0], [1,1], \ldots [10,10]$. Solving SDP (3.39) with m = 2 and m = 4 for each of these "intervals" gives the upper bounding histograms shown in Figure 3-11. For comparison, each of the plots in Figure 3-11 also shows the histogram of $P_{\rm ss}(x_{\rm A})$ values for the true steady-state probability distribution, obtained through solving the CME.

One curious feature of these plots is that for all odd x_A we have $P_{ss}(x_A)^U > 0$ while $P_{ss}(x_A) = 0$. This discrepancy, like the gaps in Figures 3-2 and 3-7, is attributable to relaxing the integrality requirement. A simple analysis of the stoichiometry of the



Figure 3-10: Bounds on a histogram for the Michaelis-Menten System at steady state



Figure 3-11: Bounds on a histogram for the Reversible Dimerization System at steady state

reaction system makes it clear that it is not actually possible to have odd counts of A molecules, given the initial condition. So one could reasonably just skip over solving SDP (3.39) for these "intervals", knowing a priori that the upper bound on probability is zero. However, we don't do that here just to be clear that this sort of parity analysis is not a built-in feature of SDP (3.39).

This parity issue aside, as with the previous example, the shape of the bounding histogram closely follows that of the true histogram.



Figure 3-12: Bounds on a histogram for the Schlögl System at steady state and relative probabilities for the true distribution. See Appendix B for an explanation of how the relative probabilities were calculated.

3.8.3 Schlögl

For the Schlögl System (3.34), we assume the same values for the rate constants and initial counts stated in Section 3.4.3 – the one exception being c_4 , which we now hold fixed at $c_4 = 3.5 \text{ s}^{-1}$. We aim to obtain bounds on the molecular count histogram for species X. Analyzing the reaction system, we see that there is no upper bound on the molecular count of X. However, we see from Figure 3-3 that with $c_4 = 3.5 \text{ s}^{-1}$, the mean $\langle X \rangle$ is around 40. Accordingly, we will construct a bounding histogram from $x_{\rm X} = 0$ to $x_{\rm X} = 99$, hoping that this range will be large enough to capture the relevant features of the distribution. In principle, we could use singleton intervals $[0,0], [1,1], \ldots, [99,99]$ to obtain the bounding histogram, much as we did for the previous two examples. However, this would require us to solve 100 SDPs. It could be done, but the relatively large molecular counts of this reaction system make it better suited to using non-singleton intervals. In particular, we will use the intervals $[0, 4], [5, 9], \dots, [95, 99]$, so that we need to solve only 20 SDPs. Solving SDP (3.39) with m = 2 for each of these intervals gives the upper bounding histogram shown in Figure 3-12. When we attempted to construct a similar bounding histogram using m = 4, we encountered numerical difficulties.


Figure 3-13: Bounds on a histogram for the Schlögl System at steady state

The upper bounding histogram we see in Figure 3-12 has a distinctly different character from those we saw in our previous two examples. In particular, it is multimodal. This is interesting because the true steady-state probability histogram for the Schlögl model with the given parameter values is also multi-modal (see Figure 3-12)

Moreover, if we increase c_4 to 5 s⁻¹, and repeat the analysis, we obtain a unimodal upper bounding histogram; and the true steady state histogram associated with this parameter value is also unimodal (see Figure 3-13). This suggests that the upper bounding histogram can serve as an indicator of multi-modality; for each peak in the true steady state distribution, there is a corresponding peak in the upper bounding histogram.

This is an intriguing idea. However, we must point out that it has its limitations. In particular, notice that the central peak of the upper bounding histogram in Figure 3-12 has no corresponding peak in true steady-state histogram. This extra peak seems to be an artifact of our bounding method, with no connection to the true distribution. Thus, while the upper bounding histogram can be used as an indicator of multi-modality, it may over-estimate the number of modes.



Figure 3-14: Bounds on a histogram for the larger reaction system in Figure 3-4 at steady state.

3.8.4 Larger Example

For the larger reaction system in Figure 3-4, we assume the same values for the rate constants and initial counts stated in Section 3.4.4 – the one exception being c_1 , which we now hold fixed at $c_1 = 1 \text{ s}^{-1}$. We aim to obtain bounds on the molecular count histogram for species A. From Figure 3-5, the mean $\langle A \rangle$ is around 24. Accordingly, we will pick as intervals $[0, 2], [3, 5], \ldots, [57, 59]$, hoping that these 20 intervals will cover a large enough range to capture the relevant features of the distribution. Solving SDP (3.39) with m = 2 for each of these intervals gives the upper bounding histogram shown in Figure 3-14. When we attempted to obtain tighter bounds by repeating the analysis with m = 4, we encountered numerical difficulties.

Because of the complexity of the reaction system, we don't know the true steadystate distribution for the count of A molecules. However, using Figure 3-14, we may speculate that it is unimodal and centered around $x_A \approx 24$.

3.9 Other Probability Bounds

The previous two sections focused on our ability to bound the probability that the molecular count of independent species *i* lies in a specific interval $[x_{\min}, x_{\max}]$. How-

ever, more generally, it is possible to bound the probability that the vector of molecular counts $\mathbf{\hat{X}}$ lies in any basic semi-algebraic set, i.e., a set of the form

$$\{\hat{\mathbf{x}} \in \mathbb{R}^{\hat{N}} : g_j(\hat{\mathbf{x}}) \ge 0, \ j = 1, \dots, K\},\tag{3.40}$$

where each $g_j(\cdot)$ for $j = 1, \ldots, K$ is a polynomial in $\hat{\mathbf{x}}$. The interval $\{\hat{\mathbf{x}} \in \mathbb{R}^{\hat{N}} : \hat{x}_i \in [x_{\min}, x_{\max}]\}$ is merely a special case of this type of set, as it can be written equivalently as $\{\hat{\mathbf{x}} \in \mathbb{R}^{\hat{N}} : \hat{x}_i - x_{\min} \ge 0, x_{\max} - \hat{x}_i \ge 0\}$. Basic semi-algebraic sets can also be used to describe polyhedra, discs, ellipses, and certain nonconvex and disconnected sets. For example, they can be used to describe the sets corresponding to the statements "the molecular counts of species A and B together do not exceed 10" or "the molecular count of species A is at least twice the molecular count of species B squared". It is not immediately obvious that bounding the probability of such events would be useful – this is largely why we chose to focus on the interval case – but it could be done.

The SDP for calculating an upper bound on the probability that the steady-state distribution assigns to an arbitrary basic semi-algebraic set is just a variation of SDP (3.39). Inspecting SDP (3.39), it is clear that the set of interest (in this case the interval $\{\hat{\mathbf{x}} \in \mathbb{R}^{\hat{N}} : \hat{x}_i \in [x_{\min}, x_{\max}]\}$) is reflected in the final LMI. To bound the probability associated with an arbitrary basic semi-algebraic set, we have to derive the LMIs corresponding to that set, and substitute them in place of the last LMI in SDP (3.39). An explanation of the derivation of these LMIs is beyond the scope of this chapter. However, Appendix B gives some insight. Further details can be found in Chapter 3 of Lasserre's book[42].

3.10 Size of the Semidefinite Programs

In this section, we tabulate the size of each of the SDPs considered in the previous examples.

The factors determining the size of the SDP include:

- the number of independent molecular species, \hat{N} ;
- the order of the highest-order reaction occuring in the system, q;
- the order of the highest-order moment whose time derivative we are setting equal to zero, m;
- and the quantity being bounded (i.e., mean, variance, or bound on probability).

Our measures of the SDP size will include:

- the order of the highest-order moment that appears as a decision variable, h;
- the number of decision variables, d;
- the number of linear constraints, *l*;
- the dimension of the largest matrix appearing in an LMI, L.

The connection between the factors determining the size and the measures h, d, l, and L, requires some explanation and is deferred to Appendix B.

3.11 Conclusion

The proposed bounding method is a fundamentally different approach to the problem of characterizing steady-state distributions in stochastic chemical kinetic systems. In this chapter, it has been applied to calculate bounds on the mean count of molecules of each species, variances in these counts, and the probability that the count lies in a given interval – leading to bounding histograms. These bounds are mathematically rigorous and based on no approximations or unprovable assumptions about the underlying probability distribution. This sets the bounding method apart from the various moment closure schemes in the literature, which provide only estimates of the quantities of interest, with no bounds on the error.

Admittedly, bounds on a quantity of interest are less useful than the value of the quantity itself. However, guaranteed bounds are arguably more useful than an

| Model | \hat{N} | q | Bounded Quantity | m | h | d | l | L |
|----------------------------|-----------|---|--|----|----|------|-----|-----|
| Michaelis- Menten | 2 | 2 | $\langle \hat{X}_i angle_{ m ss}$ | 2 | 4 | 15 | 6 | 6 |
| | | | | 4 | 6 | 28 | 15 | 10 |
| | | | $(\sigma_i^2)_{ m ss}$ | 2 | 4 | 16 | 6 | 6 |
| | | | | 4 | 6 | 29 | 15 | 10 |
| | | | $P_{\rm ss}(\hat{X}_i \in [x_{\rm min}, x_{\rm max}])$ | 2 | 4 | 30 | 6 | 6 |
| | | | | 4 | 6 | 56 | 15 | 10 |
| Reversible Dimerization | 1 | 2 | $\langle \hat{X}_i angle_{ m ss}$ | 2 | 4 | 5 | 3 | 3 |
| | | | | 4 | 6 | 7 | 5 | 4 |
| | | | $(\sigma_i^2)_{ m ss}$ | 2 | 4 | 6 | 3 | 3 |
| | | | | 4 | 6 | 8 | 5 | 4 |
| | | | $P_{\rm ss}(\hat{X}_i \in [x_{\rm min}, x_{\rm max}])$ | 2 | 4 | 10 | 3 | 3 |
| | | | | 4 | 6 | 14 | 5 | 4 |
| Schlögl | 1 | 3 | $\langle \hat{X}_i angle_{ m ss}$ | 2 | 4 | 5 | 3 | 3 |
| | | | | 6 | 8 | 9 | 7 | 5 |
| | | | | 10 | 12 | 13 | 11 | 7 |
| | | | $(\sigma_i^2)_{ m ss}$ | 2 | 4 | 6 | 3 | 3 |
| | | | | 4 | 6 | 8 | 5 | 4 |
| | | | $P_{\rm ss}(\hat{X}_i \in [x_{\min}, x_{\max}])$ | 2 | 4 | 10 | 3 | 3 |
| | | | | 4 | 6 | 14 | 5 | 4 |
| Larger Example | 7 | 2 | $\langle \hat{X}_i angle_{ m ss}$ | 2 | 4 | 330 | 36 | 36 |
| | | | | 4 | 6 | 1716 | 330 | 120 |
| | | | $(\sigma_i^2)_{ m ss}$ | 2 | 4 | 331 | 36 | 36 |
| | | | | 4 | 6 | 1717 | 330 | 120 |
| | | | $P_{\rm ss}(\hat{X}_i \in [x_{\min}, x_{\max}])$ | 2 | 4 | 660 | 36 | 36 |
| | | | | 4 | 6 | 3432 | 330 | 120 |

Table 3.1: Sizes of the SDPs solved in the examples of this chapter.

estimate with unbounded error. Furthermore, as we saw in the case of several of the mean-bounding examples, the bounds can be so tight that they effectively give the quantity of interest.

The bounding method has its limitations. First, the bounds are not always as tight as we would like for them to be. In principle, we could tighten them by increasing m (increasing the number of moments whose time derivatives we are setting equal to zero), or by selectively adding LMIs to partially enforce the integrality requirement. However, both of these strategies increase the size of the SDP that must be solved. We then confront the bounding method's second major limitation: while SDPs in general have nice theoretical properties, something about the structure of moment-based SDPs specifically seems to make them numerically ill-conditioned. This limits the size of the systems that the bounding method can handle, with current SDP solver technology. We hope that continued refinement of this technology (or perhaps development of a solver specialized for moment-based SDPs) will extend the applicability of the bounding method to ever larger systems.

In summary, we believe that the bounding method described in this chapter is an interesting theoretical addition to the wide array of tools available for analyzing stochastic chemical kinetic systems. We have demonstrated its potential by applying it to a nontrivial reaction system with over 74 billion states. However, there are some numerical issues that need to be worked out before the bounding method can be reliably applied to larger systems of practical importance.

While this chapter has been concerned with applying the bounding method to *steady-state* distributions, it can also be extended to dynamic distributions. This extension will be the subject of Chapter 4.

3.12 Implementation Details

All numerical examples in this chapter were computed on a 64-bit Dell Precision T3610 workstation with a 3.70 GHz Intel Xeon CPU. In each example, CVX [28] was used to model the SDP, using the default tolerance (i.e. "precision") settings. SeDuMi [66] was used as the underlying solver.

Chapter 4

Dynamic Bounds on Stochastic Chemical Kinetic Systems Using Semidefinite Programming

The contents of this chapter were published as a peer-reviewed paper [15] in the Journal of Chemical Physics.

4.1 Introduction

In Chapter 3, we showed that the closure problem of stochastic chemical kinetics can be partially overcome using moment-based semidefinite programs (SDPs). In particular, we showed that moment-based SDPs can be used to calculate rigorous bounds on various descriptions of the stochastic chemical kinetic system's stationary distribution(s) – for example, mean molecular counts, variances in these counts, and so on. In this chapter, we show that these ideas can be extended to the corresponding dynamic problem, calculating time-varying bounds on the same descriptions.

4.2 Mathematical Background

This section is used primarily to introduce the notation that will be used in Sections 4.3 - 4.8. Many of the ideas contained herein are similar to the corresponding section of Chapter 3. However, there are some key differences, most notably in our handling of the invariants, which foreshadow developments later in the chapter.

4.2.1 Mathematical Notation

In this chapter, the symbols \mathbb{N} , \mathbb{Z} , \mathbb{R} will be used to represent the natural numbers, integers, and real numbers, respectively. For notational convenience, we will adopt the definition of \mathbb{N} that includes 0. We will use angular brackets " $\langle \cdot \rangle$ " to represent an "expected value" or mean of a random variable. Vectors and matrices will be represented with bold symbols. The vector $\mathbf{e}_i = (0, \ldots, 1, \ldots, 0)$ is the *i*th coordinate vector, in which all components are zero, except the *i*th component, which is 1. The interpretations of all other symbols should be clear from the context in which they appear.

4.2.2 Stochastic Chemical Kinetics Notation

As in Chapter 3, we will let $N \in \mathbb{N}$ denote the number of distinct chemical species in our reaction system, and let $R \in \mathbb{N}$ denote the number of reactions. Furthermore, we will model the state of the system at time t with the random vector $\mathbf{X}(t) = (X_1(t), \ldots, X_N(t)) \in \mathbb{N}^N$, where $X_i(t) \geq 0$ is the count of molecules of species ipresent. The vector $\mathbf{s}_r \in \mathbb{Z}^N$ will contain the stoichiometric coefficients for reaction $r \in \{1, \ldots, R\}$. Concatenating these vectors together gives the stoichiometry matrix $\mathbf{S} \equiv [\mathbf{s}_1 \ldots \mathbf{s}_R] \in \mathbb{Z}^{N \times R}$. This stoichiometric matrix together with the initial state of the system $\mathbf{X}(0) \equiv \mathbf{x}_0 \in \mathbb{N}^N$ defines the set of all reachable states:

$$\mathcal{X} \equiv \left\{ \mathbf{x} \in \mathbb{N}^N : \begin{array}{c} \mathbf{x} = \mathbf{x}_0 + \mathbf{S}\mathbf{y}, \\ \mathbf{x} \ge \mathbf{0}, \quad \mathbf{y} \in \mathbb{N}^R \end{array} \right\}.$$
(4.1)

4.2.3 Invariants and Independent Species

Let the vectors $\{\mathbf{b}_1, \ldots, \mathbf{b}_L\} \subset \mathbb{R}^N$ form a basis for the left null space of the stoichiometry matrix **S**. Each vector \mathbf{b}_j corresponds to an *invariant* of the reaction system[20]. In particular,

$$\mathbf{b}_{j}^{\mathrm{T}}\mathbf{X}(t) = f_{j}, \quad \forall j \in \{1, \dots, L\}, \quad \forall t \ge 0,$$

$$(4.2)$$

where each $f_j \in \mathbb{R}$ is a constant which we will call the *value* of the *j*th invariant. This is a small but significant deviation from Chapter 3, where Equation (4.2) was written not in terms of invariant values but in terms of the initial state $\mathbf{X}(0)$. In what follows, we will assume only that the invariant values $\{f_1, \ldots, f_L\}$ are known. Precise knowledge of the initial state $\mathbf{X}(0)$ is unnecessary. Of course, knowledge of $\mathbf{X}(0)$ might be used to calculate the invariant values $\{f_1, \ldots, f_L\}$, via Equation (4.2). However, the definition of the invariants does not depend explicitly on $\mathbf{X}(0)$. This has some interesting implications regarding uncertainty in the initial state, which will be explored further in Section 4.8.

If we set $\mathbf{f} \equiv (f_1, \ldots, f_L)$, and

$$\mathbf{B} \equiv \begin{bmatrix} \mathbf{b}_{1}^{\mathrm{T}} \\ \vdots \\ \mathbf{b}_{L}^{\mathrm{T}} \end{bmatrix} \in \mathbb{R}^{L \times N}, \tag{4.3}$$

then Equation (4.2) can be expressed concisely as

$$\mathbf{BX}(t) = \mathbf{f}, \quad \forall t \ge 0. \tag{4.4}$$

These equations imply that the set of reachable states \mathcal{X} is contained in an affine subspace, i.e., that $\mathcal{X} \subset \{\mathbf{x} \in \mathbb{R}^n : \mathbf{B}\mathbf{x} = \mathbf{f}\}$. Furthermore, they imply that not all molecular counts X_1, \ldots, X_N can vary independently. In particular, following the same basic idea of Chapter 3, we can write

$$\tilde{\mathbf{X}}(t) = \tilde{\mathbf{B}}^{-1}\mathbf{f} - \tilde{\mathbf{B}}^{-1}\hat{\mathbf{B}}\hat{\mathbf{X}}(t), \quad \forall t \ge 0,$$
(4.5)

where $\tilde{\mathbf{B}} \in \mathbb{R}^{L \times L}$ is a matrix obtained by concatenating L linearly independent columns of \mathbf{B} , and the matrix $\hat{\mathbf{B}} \in \mathbb{R}^{L \times \hat{N}}$ is obtained by concatenating the remaining $N - L \equiv \hat{N}$ columns. Furthermore, the vector $\tilde{\mathbf{X}}(t) \in \mathbb{N}^L$ consists of the components of $\mathbf{X}(t)$ corresponding to the columns of $\tilde{\mathbf{B}}$, and, similarly, $\hat{\mathbf{X}}(t) \in \mathbb{N}^{\hat{N}}$ consists of the components of $\mathbf{X}(t)$ corresponding to the columns of $\hat{\mathbf{B}}$. Equation (4.5) allows us to treat the components represented in the vector $\hat{\mathbf{X}}$ as the "independent species" and the components represented in the vector $\tilde{\mathbf{X}}$ as the "dependent species".

4.2.4 A Reduced State Space

Since the molecular counts of the independent species determine the molecular counts of the dependent species through Equation (4.5), each full-dimensional reachable state $\mathbf{x} \in \mathcal{X} \subset \mathbb{N}^N$ is equivalent to a *reduced reachable state*, $\hat{\mathbf{x}} \in \mathbb{N}^{\hat{N}}$, which specifies only the molecular counts of the independent species. We will represent the set of all these reduced reachable states as $\hat{\mathcal{X}} \subset \mathbb{N}^{\hat{N}}$.

We know that the molecular counts of the independent species must be nonnegative, so for any $\hat{\mathbf{x}} \in \hat{\mathcal{X}}$, we must have $\hat{\mathbf{x}} \ge \mathbf{0}$. Furthermore, we know that the molecular counts of the dependent species must be nonnegative. By Equation (4.5), this implies $\tilde{\mathbf{B}}^{-1}\mathbf{f} - \tilde{\mathbf{B}}^{-1}\hat{\mathbf{B}}\hat{\mathbf{x}} \ge \mathbf{0}$. It follows that the set of reduced reachable states $\hat{\mathcal{X}}$ must be contained in the following polyhedral set:

$$\bar{\mathcal{X}} \equiv \left\{ \hat{\mathbf{x}} \in \mathbb{R}^{\hat{N}} : \begin{array}{c} \hat{\mathbf{x}} \ge \mathbf{0}, \\ \tilde{\mathbf{B}}^{-1}\mathbf{f} - \tilde{\mathbf{B}}^{-1}\hat{\mathbf{B}}\hat{\mathbf{x}} \ge \mathbf{0} \end{array} \right\}.$$
(4.6)

For the sake of brevity, in what follows, we will often loosely refer to the "state" of the system, meaning the "reduced reachable state", $\hat{\mathbf{x}} \in \hat{\mathcal{X}}$. If we are instead talking about the full-dimensional state $\mathbf{x} \in \mathcal{X}$, this will be clear from the context.

4.2.5 The Chemical Master Equation

As described in Chapter 3, we assign a probability $Pr(\hat{\mathbf{X}}(t) = \hat{\mathbf{x}}, t) \equiv P(\hat{\mathbf{x}}, t)$ to each of the reachable states $\hat{\mathbf{x}} \in \hat{\mathcal{X}}$. Taken together, this collection of probability values constitutes a probability distribution $P(\cdot, t)$ supported on the set of reachable states $\hat{\mathcal{X}}$. This distribution changes over time according to the chemical master equation (CME):

$$\frac{dP}{dt}(\hat{\mathbf{x}},t) = \sum_{r=1}^{R} [P(\hat{\mathbf{x}} - \hat{\mathbf{s}}_r, t) a_r(\hat{\mathbf{x}} - \hat{\mathbf{s}}_r) - P(\hat{\mathbf{x}}, t) a_r(\hat{\mathbf{x}})],$$

$$\forall \hat{\mathbf{x}} \in \hat{\mathcal{X}},$$
(4.7)

where $\hat{\mathbf{s}}_r \in \mathbb{Z}^{\hat{N}}$ is the vector obtained by selecting only the components of the vector $\mathbf{s}_r \in \mathbb{Z}^N$ corresponding to the independent species, and a_r is the "propensity function" of reaction r (see Higham[31] and Gillespie[26] for details). In the general case, the propensity functions depend on all molecular counts, and are thus often written as a function of the full-dimensional vector $\mathbf{x} \in \mathbb{N}^N$. However, since the dependent species can be expressed in terms of the independent species via $\tilde{\mathbf{x}} = \tilde{\mathbf{B}}^{-1}\mathbf{f} - \tilde{\mathbf{B}}^{-1}\hat{\mathbf{B}}\hat{\mathbf{x}}$, the propensity functions can be considered as functions of the independent species $\hat{\mathbf{x}} \in \mathbb{N}^{\hat{N}}$ alone. In this chapter, we consider only the case of mass-action kinetics, which implies that the propensity functions are polynomials of $\hat{\mathbf{x}}$. In principle, our method can handle polynomials of arbitrarily high order (and even rational functions, if we extend the results of Kuntz et al[39]), but we focus on the physically important special case of second-order polynomials, assuming reactions requiring at most bimolecular collisions.

If we specify an initial probability distribution $P(\cdot, 0)$, the CME determines all future probability distributions $P(\cdot, t)$ for t > 0. Often this initial distribution is assumed to be a Dirac distribution, $P(\cdot, 0) = \delta_{\hat{\mathbf{x}}_0}$, where all of the probability is concentrated on a single state $\hat{\mathbf{x}}_0 \in \hat{\mathcal{X}}$. However, in principle, the initial distribution could be supported on any subset of $\hat{\mathcal{X}}$.

As described by several authors [49, 72], the CME can also be expressed as a linear

time-invariant ordinary differential equation (LTI ODE):

$$\frac{d\mathbf{p}}{dt}(t) = \mathbf{G}\mathbf{p}(t),\tag{4.8}$$

where **p** is a vector specifying the probability for each state $\hat{\mathbf{x}} \in \hat{\mathcal{X}}$, and **G** is the time-invariant "infinitesimal generator" matrix describing how probability flows from one state to another. Conceptually simple as this equation is, it is usually too large to compute a numerical solution.

4.2.6 Moments in Stochastic Chemical Kinetics

The probability distribution $P(\cdot, t)$ can be characterized by its moments, defined as

$$\mu_{\mathbf{j}}(t) \equiv \sum_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}} \hat{\mathbf{x}}^{\mathbf{j}} P(\hat{\mathbf{x}}, t), \tag{4.9}$$

for any multi-index $\mathbf{j} = (j_1, \ldots, j_{\hat{N}}) \in \mathbb{N}^{\hat{N}}$, where the sum is over the set $\hat{\mathcal{X}}$ of all reachable states, and $\hat{\mathbf{x}}^{\mathbf{j}} = \prod_{k=1}^{\hat{N}} \hat{x}_k^{j_k}$ is a monomial. The *order* of the moment $\mu_{\mathbf{j}}$ is defined as the sum $|\mathbf{j}| \equiv \sum_{k=1}^{\hat{N}} j_k$. Notice that $\mu_{\mathbf{0}}(t) = 1$ for all times t.

The first-order moment $\mu_{\mathbf{e}_i}(t)$ is the mean molecular count for independent species $i \in \{1, \dots, \hat{N}\}$ at time t:

$$\mu_{\mathbf{e}_{i}}(t) \equiv \sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}} \hat{\mathbf{x}}^{\mathbf{e}_{i}} P(\hat{\mathbf{x}}, t) = \sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}} \hat{x}_{i} P(\hat{\mathbf{x}}, t) = \langle \hat{X}_{i}(t) \rangle.$$
(4.10)

Furthermore, if we let $\beta_{k,j}$ denote the element in the *k*th row and *j*th column of the matrix $\tilde{\mathbf{B}}^{-1}\hat{\mathbf{B}}$, and α_k equal the *k*th component of the vector $\tilde{\mathbf{B}}^{-1}\mathbf{f}$, then Equation (4.5) allows us to express the mean molecular count for each dependent species $k \in$

 $\{1,\ldots,L\}$:

$$\begin{split} \langle \tilde{X}_{k}(t) \rangle &= \sum_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}} \tilde{x}_{k} P(\hat{\mathbf{x}}, t), \\ &= \sum_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}} \mathbf{e}_{k}^{\mathrm{T}} \left(\tilde{\mathbf{B}}^{-1} \mathbf{f} - \tilde{\mathbf{B}}^{-1} \hat{\mathbf{B}} \hat{\mathbf{x}} \right) P(\hat{\mathbf{x}}, t), \\ &= \sum_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}} \left(\alpha_{k} - \sum_{j=1}^{\hat{N}} \beta_{k,j} \hat{x}_{j} \right) P(\hat{\mathbf{x}}, t), \\ &= \alpha_{k} - \sum_{j=1}^{\hat{N}} \beta_{k,j} \mu_{\mathbf{e}_{j}}(t). \end{split}$$
(4.11)

As for the second-order moments, we see that $\mu_{\mathbf{e}_i}(t)$ and $\mu_{2\mathbf{e}_i}(t)$ can be used together to compute the variance in the count of molecules of independent species *i* at time *t*:

$$\sigma_i^2(t) \equiv \langle \hat{X}_i^2(t) \rangle - \langle \hat{X}_i(t) \rangle^2 = \mu_{2\mathbf{e}_i}(t) - \mu_{\mathbf{e}_i}^2(t).$$
(4.12)

4.2.7 The Closure Problem

As described by several authors[64, 65, 25], the CME can be used to derive a system of linear ordinary differential equations describing how the moments of the distribution $P(\cdot, t)$ change over time. For reaction systems containing at most unimolecular (i.e., first-order) reactions, this equation can be written as

$$\frac{d\boldsymbol{\mu}_L}{dt}(t) = \mathbf{A}_L \boldsymbol{\mu}_L(t) \tag{4.13}$$

where $\boldsymbol{\mu}_L(t)$ is a vector of "low-order" moments up to some arbitrary order $m \in \mathbb{N}$, and \mathbf{A}_L is a constant square matrix of appropriate dimension. This simple ODE poses no problem. However, for systems in which the highest order reaction is order q > 1, the ODE becomes

$$\frac{d\boldsymbol{\mu}_L}{dt}(t) = \mathbf{A}_L \boldsymbol{\mu}_L(t) + \mathbf{A}_H \boldsymbol{\mu}_H(t), \qquad (4.14)$$

where $\boldsymbol{\mu}_{H}(t)$ is a vector of "high-order" moments, order m + 1 through $m + q - 1 \equiv M$; and \mathbf{A}_{H} is an appropriately sized constant matrix, usually non-square. In the physically important special case where the reaction system contains at most

bimolecular reactions (i.e., q = 2), the vector $\boldsymbol{\mu}_H(t)$ contains only moments of order m + 1 = M. What Equation (4.14) expresses is that the time derivatives of the low-order moments depend not only on the values of the low-order moments but also on the values of the high-order moments – which are themselves changing in time. This is the case no matter what value of m we pick as our arbitrary cut-off of what to consider a "low-order" moment. It is unclear how to solve such a dynamic system. This is the infamous "closure problem".

4.3 Bounds on Dynamic Systems

In this section, we describe the main contribution of this chapter: a method for calculating time-varying bounds on various quantities for stochastic chemical kinetic systems. This method is an extension of our previous work on calculating steady-state bounds for such systems[14], and some similarities will be apparent. However, considering dynamics adds an element of complexity not found in our Chapter 3.

4.3.1 The Paradigm

Suppose that we have a generic stochastic chemical kinetic system, characterized by a stoichiometry matrix $\mathbf{S} \in \mathbb{Z}^{N \times R}$ and a vector of rate constants $\mathbf{c} \in \mathbb{R}^{R}$. Assume that there is at least one reaction with order greater than one, so that this system exhibits the closure problem when subjected to a moment analysis. Suppose that we have analyzed \mathbf{S} to construct an invariant matrix $\mathbf{B} \in \mathbb{R}^{L \times N}$, as described in Section 4.2.3, and that we know the associated invariant values $\mathbf{f} \in \mathbb{R}^{L}$. Suppose further that have identified the $\hat{N} = N - L$ chemical species we wish to treat as independent and constructed the matrices $\hat{\mathbf{B}} \in \mathbb{R}^{L \times \hat{N}}$ and $\tilde{\mathbf{B}} \in \mathbb{R}^{L \times L}$. Finally, suppose that we have chosen a value of $m \in \mathbb{N}$ and constructed the matrices \mathbf{A}_{L} and \mathbf{A}_{H} described in Section 4.2.7. We are interested in analyzing the properties of the probability distribution describing the stochastic chemical kinetic system at a particular time T.

Consider the problem of bounding $\langle \hat{X}_i(T) \rangle$, the mean count of molecules of independent species *i* at time *T*. What we'd like to do is calculate two numbers $\langle \hat{X}_i(T) \rangle^L$

and $\langle \hat{X}_i(T) \rangle^U$ such that

$$\langle \hat{X}_i(T) \rangle^L \le \langle \hat{X}_i(T) \rangle \le \langle \hat{X}_i(T) \rangle^U$$
(4.15)

is guaranteed.

First, let us focus on the upper bound, $\langle \hat{X}_i(T) \rangle^U$. To calculate this bound, we will write down several mathematical conditions that the moments $\boldsymbol{\mu}(T)$ of the probability distribution $P(\cdot, T)$ must necessarily satisfy. We will then optimize over all of vectors $\tilde{\boldsymbol{\mu}}$ that satisfy these necessary conditions, searching for that vector which maximizes $\tilde{\mu}_{\mathbf{e}_i}(T)$, the expression corresponding to the mean molecular count of independent species *i* (see Equation (4.10)). This optimization can be stated slightly more formally, but still rather abstractly, as

$$\langle \hat{X}_i(T) \rangle^U \equiv \max_{\tilde{\mu}(T)} \quad \tilde{\mu}_{\mathbf{e}_i}(T)$$
s.t. $\tilde{\mu}(T)$ satisfies necessary (4.16)
moment conditions at time T .

Note that we are making a distinction between the true moment vector $\boldsymbol{\mu}(T)$ at time T, and the decision variable $\tilde{\boldsymbol{\mu}}(T)$, which is a proxy for $\boldsymbol{\mu}(T)$. By construction, the true moment vector $\boldsymbol{\mu}(T)$ is a feasible point for this optimization problem. This implies that the problem's optimal value $\langle \hat{X}_i(T) \rangle^U$ is guaranteed to be an upper bound on $\mu_{\mathbf{e}_i}(T) = \langle \hat{X}_i(T) \rangle$, the true mean molecular count of species *i* at time *T*.

Now, whether or not we can actually solve this problem remains to be seen. The sole point of this section is a theoretical one: *if* we can solve the abstract optimization problem above, we would have our desired upper bound.

Notice that the above reasoning is valid whether $\tilde{\boldsymbol{\mu}}(T)$ and $\boldsymbol{\mu}(T)$ are considered to be infinite sequences or vectors containing only finitely many moments. However, for practical computations, we must work with finite vectors. So, going forward, we will specify that $\tilde{\boldsymbol{\mu}}(T)$ contains only those moments up to order $2n \in \mathbb{N}$, where $n = \lceil \frac{M}{2} \rceil$. The reason for this choice of n is explained in detail in Chapter 3. Roughly speaking, it ensures that we are explicitly representing enough moments to express the dynamics described by Equation (4.14).

4.3.2 Necessary Moment Conditions

What exactly are the necessary moment conditions appearing in Problem (4.16)? First, we must have that the total probability is equal to one:

$$\mu_0(T) = 1, \tag{4.17}$$

Second, as explained in Chapter 3, the fact that the distribution $P(\cdot, T)$ is supported on the set $\hat{\mathcal{X}} \subset \bar{\mathcal{X}}$ implies that its moments $\boldsymbol{\mu}(T)$ must satisfy a set of linear matrix inequalities (LMIs):

$$\mathbf{M}_{n}^{\mathbf{0}}(\boldsymbol{\mu}(T)) \succeq \mathbf{0}, \tag{4.18}$$

$$\mathbf{M}_{n-1}^{\mathbf{e}_j}(\boldsymbol{\mu}(T)) \succeq \mathbf{0}, \quad \forall j \in \{1, \dots, \hat{N}\},$$
(4.19)

$$\alpha_{k} \mathbf{M}_{n-1}^{\mathbf{0}}(\boldsymbol{\mu}(T)) - \sum_{j=1}^{\hat{N}} \beta_{k,j} \mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\boldsymbol{\mu}(T)) \succeq \mathbf{0},$$

$$\forall k \in \{1, \dots, L\}.$$
 (4.20)

The exact definitions of the matrices $\mathbf{M}_{n}^{\mathbf{0}}(\boldsymbol{\mu}(T))$, $\mathbf{M}_{n-1}^{\mathbf{0}}(\boldsymbol{\mu}(T))$, and $\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\boldsymbol{\mu}(T))$ can be found in Appendix B, but are unimportant for the present discussion. What is important is that these matrices are symmetric and linear with respect to their arguments. Each LMI simply asserts that the matrices on the left-hand side of the " \succeq " must be positive semidefinite (i.e., have all nonnegative eigenvalues).

The set of vectors satisfying LMIs (4.18)-(4.20) is a mathematical cone. To simplify the notation in what follows, will represent this cone concisely as $C_n(\boldsymbol{\alpha}, \boldsymbol{\beta})$. Thus,

$$\boldsymbol{\mu}(T) \in C_n(\boldsymbol{\alpha}, \boldsymbol{\beta}) \tag{4.21}$$

is equivalent to LMIs (4.18)-(4.20).

Necessary Conditions (4.17) and (4.21) are notably lacking any information about

the dynamics of the system. To obtain necessary conditions implied by the dynamics, we make use of Equation (4.14), which holds for all times t. Suppose that we pick an arbitrary $\rho \in \mathbb{R}$, multiply both sides of Equation (4.14) by $e^{\rho(T-t)}$, and then integrate from t = 0 to t = T:

$$\int_0^T e^{\rho(T-t)} \frac{d\boldsymbol{\mu}_L}{dt}(t) dt$$

$$= \int_0^T e^{\rho(T-t)} (\mathbf{A}_L \boldsymbol{\mu}_L(t) + \mathbf{A}_H \boldsymbol{\mu}_H(t)) dt.$$
(4.22)

Applying integration by parts to the left-hand side, we obtain

$$\int_{0}^{T} e^{\rho(T-t)} \frac{d\boldsymbol{\mu}_{L}}{dt}(t) dt$$

= $e^{\rho(T-t)} \boldsymbol{\mu}_{L}(t) |_{0}^{T} - \int_{0}^{T} (-\rho) e^{\rho(T-t)} \boldsymbol{\mu}_{L}(t) dt,$ (4.23)
= $\boldsymbol{\mu}_{L}(T) - e^{\rho T} \boldsymbol{\mu}_{L}(0) + \rho \int_{0}^{T} e^{\rho(T-t)} \boldsymbol{\mu}_{L}(t) dt.$

We presume that the initial values of the low-order moments $\boldsymbol{\mu}_L(0)$ can be easily computed from the initial distribution $P(\cdot, 0)$ via Equation (4.9). This is true, for example, if the initial molecular count is known exactly – which corresponds to an initial probability distribution $P(\cdot, 0)$ where all the probability is concentrated on a single state $\hat{\mathbf{x}}_0$, i.e., the Dirac distribution $\delta_{\hat{\mathbf{x}}_0}$. However, it may also be the case that we *don't* know the initial molecular count exactly. In this case, our initial probability distribution $P(\cdot, 0)$ will be supported on several reachable states $\hat{\mathbf{x}} \in \hat{\mathcal{X}}$. Our method can handle this situation, as well, as long as we can compute the moments $\boldsymbol{\mu}_L(0)$ (see Section 4.8).

For the right-hand side, we can make use of the fact that the integral is a linear operator to obtain

$$\mathbf{A}_L \int_0^T e^{\rho(T-t)} \boldsymbol{\mu}_L(t) dt + \mathbf{A}_H \int_0^T e^{\rho(T-t)} \boldsymbol{\mu}_H(t) dt.$$
(4.24)

If we define the generalized moments

$$z_{\mathbf{j}}^{(\rho)} \equiv \int_0^T e^{\rho(T-t)} \mu_{\mathbf{j}}(t) dt, \quad \forall \mathbf{j} \in \mathbb{N}^{\hat{N}},$$
(4.25)

we can express Equation (4.22) concisely as

$$\boldsymbol{\mu}_{L}(T) - e^{\rho T} \boldsymbol{\mu}_{L}(0) + \rho \mathbf{z}_{L}^{(\rho)} = \mathbf{A}_{L} \mathbf{z}_{L}^{(\rho)} + \mathbf{A}_{H} \mathbf{z}_{H}^{(\rho)}.$$
(4.26)

Rearranging, we obtain the following necessary condition on $\mu(T)$:

$$\boldsymbol{\mu}_{L}(T) - e^{\rho T} \boldsymbol{\mu}_{L}(0) = (\mathbf{A}_{L} - \rho \mathbf{I}) \mathbf{z}_{L}^{(\rho)} + \mathbf{A}_{H} \mathbf{z}_{H}^{(\rho)}.$$
(4.27)

Now, recall that we are deriving these necessary conditions with the intent of using them as constraints in our abstract Problem (4.16). In particular, we would like to write each necessary condition in terms of the decision variable proxy $\tilde{\mu}_L(T)$. For Necessary Conditions (4.17) and (4.21), this is straightforward. They translate into the constraints

$$\tilde{\mu}_0(T) = 1 \tag{4.28}$$

and

$$\tilde{\boldsymbol{\mu}}(T) \in C_n(\boldsymbol{\alpha}, \boldsymbol{\beta}). \tag{4.29}$$

For Necessary Condition (4.27), we are tempted to write the constraint

$$\tilde{\boldsymbol{\mu}}_{L}(T) - e^{\rho T} \boldsymbol{\mu}_{L}(0) = (\mathbf{A}_{L} - \rho \mathbf{I}) \mathbf{z}_{L}^{(\rho)} + \mathbf{A}_{H} \mathbf{z}_{H}^{(\rho)}.$$
(4.30)

However, there is a problem here: while the vectors $\mathbf{z}_{L}^{(\rho)}$ and $\mathbf{z}_{H}^{(\rho)}$ are well-defined, their values are unknown. This suggests that we should replace them with additional proxies $\tilde{\mathbf{z}}_{L}^{(\rho)}$ and $\tilde{\mathbf{z}}_{H}^{(\rho)}$, which will also be considered as decision variables in Problem (4.16). Doing so allows us to write the following constraint, which, importantly, is linear in the decision variables:

$$\tilde{\boldsymbol{\mu}}_{L}(T) - e^{\rho T} \boldsymbol{\mu}_{L}(0) = (\mathbf{A}_{L} - \rho \mathbf{I}) \tilde{\mathbf{z}}_{L}^{(\rho)} + \mathbf{A}_{H} \tilde{\mathbf{z}}_{H}^{(\rho)}.$$
(4.31)

4.3.3 Necessary Conditions on the Generalized Moments

Now, by itself, Equation (4.31) isn't very useful as a constraint on $\tilde{\boldsymbol{\mu}}(T)$, because it is in terms of the unknown vector $\tilde{\mathbf{z}}^{(\rho)} \equiv (\tilde{\mathbf{z}}_L^{(\rho)}, \tilde{\mathbf{z}}_H^{(\rho)})$. It tells us only that $\tilde{\boldsymbol{\mu}}_L(T) - e^{\rho T} \boldsymbol{\mu}_L(0)$ must be contained in the column space of the matrix $[(\mathbf{A}_L - \rho \mathbf{I}) \quad \mathbf{A}_H]$. However, if we can constrain the set of possible $\tilde{\mathbf{z}}^{(\rho)}$ values, Equation (4.31) becomes more useful to us. Our strategy for constraining the possible values of $\tilde{\mathbf{z}}^{(\rho)}$ will be analogous to what we've done for $\tilde{\boldsymbol{\mu}}(T)$. We will first derive necessary conditions on the generalized moments $\mathbf{z}^{(\rho)}$ and then translate them into constraints on $\tilde{\mathbf{z}}^{(\rho)}$.

The first constraint on the $\tilde{\mathbf{z}}^{(\rho)}$ values is derived from the fact that $\mu_0(t) = 1$ for all times $t \in [0, T]$. This implies

$$z_{0}^{(\rho)} = \int_{0}^{T} e^{\rho(T-t)} \mu_{0}(t) dt$$

=
$$\int_{0}^{T} e^{\rho(T-t)} dt$$

=
$$\begin{cases} T & \text{if } \rho = 0, \\ \frac{e^{\rho T} - 1}{\rho} & \text{otherwise.} \end{cases}$$
 (4.32)

which suggests the constraint

$$\tilde{z}_{\mathbf{0}}^{(\rho)} = \begin{cases} T & \text{if } \rho = 0, \\ \frac{e^{\rho T} - 1}{\rho} & \text{otherwise.} \end{cases}$$
(4.33)

The second set of constraints is derived from LMIs (4.18) - (4.20). Since these LMIs are derived solely from the fact that the unknown probability distribution is supported on $\hat{\mathcal{X}} \subset \bar{\mathcal{X}}$, they hold not just at time *T*, but also for all times $t \in [0, T]$. For example, we have

$$\mathbf{M}_{n}^{\mathbf{0}}(\boldsymbol{\mu}(t)) \succeq \mathbf{0}, \quad \forall t \in [0, T].$$

$$(4.34)$$

Multiplying both sides of the LMI by the nonnegative factor $e^{\rho(T-t)}$ and integrating

over [0, T] maintains the LMI:

$$\int_0^T e^{\rho(T-t)} \mathbf{M}_n^{\mathbf{0}}(\boldsymbol{\mu}(t)) dt \succeq \mathbf{0}.$$
(4.35)

Furthermore, because the integral is a linear operator, and because $\mathbf{M}_n^{\mathbf{0}}(\cdot)$ is a linear function of its argument, we can bring the integral inside:

$$\mathbf{M}_{n}^{\mathbf{0}}\left(\int_{0}^{T}e^{\rho(T-t)}\boldsymbol{\mu}(t)dt\right) = \mathbf{M}_{n}^{\mathbf{0}}\left(\mathbf{z}^{(\rho)}\right) \succeq \mathbf{0}.$$
(4.36)

Following similar reasoning, we can show that

$$\mathbf{M}_{n-1}^{\mathbf{e}_j}(\boldsymbol{z}^{(\rho)}) \succeq \mathbf{0}, \quad \forall j \in \{1, \dots, \hat{N}\},$$
(4.37)

$$\alpha_{k}\mathbf{M}_{n-1}^{\mathbf{0}}(\boldsymbol{z}^{(\rho)}) - \sum_{j=1}^{\hat{N}} \beta_{k,j}\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\boldsymbol{z}^{(\rho)}) \succeq \mathbf{0},$$

$$\forall k \in \{1, \dots, L\}.$$
(4.38)

LMIs (4.36) - (4.38) can be written concisely as

$$\boldsymbol{z}^{(\rho)} \in C_n(\boldsymbol{\alpha}, \boldsymbol{\beta}). \tag{4.39}$$

We have shown that membership in the cone $C_n(\boldsymbol{\alpha}, \boldsymbol{\beta})$ is a necessary condition for the vector $\boldsymbol{z}^{(\rho)}$. Accordingly, we will enforce this membership as a constraint on its decision variable proxy $\tilde{\boldsymbol{z}}^{(\rho)}$:

$$\tilde{\boldsymbol{z}}^{(\rho)} \in C_n(\boldsymbol{\alpha}, \boldsymbol{\beta}).$$
 (4.40)

Recall that our choice of $\rho \in \mathbb{R}$ was arbitrary. It follows that Necessary Conditions (4.27), (4.32), and (4.39) can be written for any $\rho \in \mathbb{R}$. In fact, they hold for each ρ in any subset $\mathcal{R} \subset \mathbb{R}$. Furthermore, we can write constraints (4.31), (4.33), and (4.40) for each ρ in any subset $\mathcal{R} \subset \mathbb{R}$.

4.3.4 A Semidefinite Program

If we use constraints (4.28), (4.29), (4.31), (4.33), and (4.40) in place of the abstract statement " $\tilde{\mu}(T)$ satisfies necessary moment conditions at time T", we obtain Problem (4.41):

$$\langle \hat{X}_{i}(T) \rangle^{U} = \max_{\substack{\boldsymbol{\mu}(T),\\ \boldsymbol{\tilde{z}}^{(\rho)}, \forall \rho \in \mathcal{R}}} \quad \tilde{\mu}_{\mathbf{e}_{i}}(T)$$
s.t. $\tilde{\mu}_{\mathbf{0}}(T) = 1,$

$$\qquad \qquad \tilde{\boldsymbol{\mu}}(T) \in C_{n}(\boldsymbol{\alpha}, \boldsymbol{\beta}), \quad \forall \rho \in \mathcal{R},$$

$$\qquad \qquad \qquad \tilde{\mathbf{z}}^{(\rho)} \in C_{n}(\boldsymbol{\alpha}, \boldsymbol{\beta}), \quad \forall \rho \in \mathcal{R},$$
Equation (4.31) holds, $\forall \rho \in \mathcal{R},$
Equation (4.33)holds, $\forall \rho \in \mathcal{R}.$

Note that the vectors $\tilde{\mathbf{z}}^{(\rho)}$ for all $\rho \in \mathcal{R}$ are decision variables in addition to the vector $\tilde{\boldsymbol{\mu}}(T)$. As with the vector $\tilde{\boldsymbol{\mu}}(T)$, it is only necessary for these vectors to contain moments up through order 2n, where $n \equiv \lceil \frac{M}{2} \rceil$.

With its linear objective function, linear equations, and LMIs, Problem (4.41) is a special type of optimization problem called a Semidefinite Program (SDP). As with all SDPs, Problem (4.41) is convex. Thus, at least in theory, we should be able to solve it efficiently[71]. Doing so, we obtain the desired upper bound, $\langle \hat{X}_i(T) \rangle^U$. Solving the corresponding minimization problem, we obtain the lower bound, $\langle \hat{X}_i(T) \rangle^L$.

4.3.5 Irrelevance of Intermediate Time Points

One intriguing property of SDP (4.41) and its minimization counterpart is that they can give us bounds on the mean molecular count of species i at time t = T using only knowledge of

- the reaction invariants (encoded in α and β) and
- the distribution at time t = 0 (encoded in the moments $\mu_L(0)$ in Equation (4.31)).

In particular, to obtain the bounds at time t = T, it is not necessary to do any calculations concerned with any intermediate time points, $t \in (0, T)$. Moreover, this is true for any value of $T \in (0, +\infty)$, no matter how large. We will revisit this idea in Section 4.9.

4.3.6 Inspiration from Previous Work

The inspiration for the bounding method described in the preceding sections comes from a paper by Bertsimas and Caramanis[3], in which moment-based SDPs are used to bound the solutions of linear partial differential equations (PDEs). The central idea of their method is to view the solution $u(\cdot)$ of the PDE as a distribution over the problem domain Ω . Taking this view, they define the *full moments*

$$m_{\mathbf{j}} \equiv \int_{\Omega} \mathbf{x}^{\mathbf{j}} u(\mathbf{x}) \tag{4.42}$$

and boundary moments

$$b_{\mathbf{j}} \equiv \int_{\partial\Omega} \mathbf{x}^{\mathbf{j}} u(\mathbf{x}) \tag{4.43}$$

of the distribution, where $\partial\Omega$ is some portion of the boundary. Starting from the linear PDE and the associated boundary conditions, they derive linear equations that these moments must satisfy. Furthermore, they derive LMIs that the moments must satisfy, simply by virtue of being moments of a distribution supported on Ω . They then solve an SDP to optimize over all vectors (**m**, **b**) which satisfy these necessary conditions, searching for that vector which maximizes or minimizes some moment of interest.

Clearly, this is thematically similar to the bounding method we have proposed for stochastic chemical kinetic systems. We now elaborate on this connection. In considering the problem of stochastic chemical kinetics, we naturally focus on $P(\cdot, t)$ as a probability distribution over the reachable states $\hat{\mathcal{X}}$ for each time $t \in [0, T]$. However, we can also think of the function $P(\cdot, \cdot)$ as a generalized distribution over both state space and time – that is, a distribution supported on the set $\Omega = \hat{\mathcal{X}} \times$ [0, T]. This $P(\cdot, \cdot)$ is directly analogous to the function $u(\cdot)$ above. Furthermore, the moments $\mu_{\mathbf{j}}(0)$ and $\mu_{\mathbf{j}}(T)$ are analogous to the "boundary moments", as they are associated with the boundaries of Ω corresponding to t = 0 and t = T. Finally, the quantities $z_{\mathbf{j}}^{(\rho)}$ are analogous to the "full moments" above.

This last analogy may not be so obvious, but it becomes clearer if we expand Equation (4.25) using Equation (4.9). Doing so, we see that for any $\mathbf{j} \in \mathbb{N}^{\hat{N}}$,

$$z_{\mathbf{j}}^{(\rho)} = \int_0^T \sum_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}} e^{\rho(T-t)} \hat{\mathbf{x}}^{\mathbf{j}} P(\hat{\mathbf{x}}, t) dt, \qquad (4.44)$$

which can be written more abstractly, closer to Bertsimas and Caramanis's notation, as

$$z_{\mathbf{j}}^{(\rho)} = \int_{\Omega} e^{\rho(T-t)} \hat{\mathbf{x}}^{\mathbf{j}} P(\hat{\mathbf{x}}, t), \qquad (4.45)$$

where, again, $\Omega = \hat{\mathcal{X}} \times [0, T]$. When the equation for $z_{\mathbf{j}}^{(\rho)}$ is written in this form, the analogy with Equation (4.42) is obvious.

The reader might protest that a closer analogy to Equation (4.42) would be

$$z_{\mathbf{j}}^{(\rho)} = \int_{\Omega} t^{\rho} \hat{\mathbf{x}}^{\mathbf{j}} P(\hat{\mathbf{x}}, t), \qquad (4.46)$$

and we agree. Our departure from the strict analogy is deliberate. As Bertsimas and Caramanis point out, while moments are classically defined in terms of monomials, we are free to define them in terms of other basis functions which may be better suited to the problem at hand. This is exactly what we have done in our definition of $z_{\mathbf{j}}^{(\rho)}$. Recall that the CME (4.8) is a linear time-invariant ODE:

$$\frac{d\mathbf{p}}{dt}(t) = \mathbf{G}\mathbf{p}(t).$$

Assuming that the number of reachable states $|\hat{\mathcal{X}}|$ is finite, and assuming that **G** has $|\hat{\mathcal{X}}|$ distinct eigenvalues $\{\lambda_j\}_{j=1}^{|\hat{\mathcal{X}}|}$, the solution to this system can be written as

$$\mathbf{p}(t) = \sum_{j=1}^{|\hat{\mathcal{X}}|} a_j e^{\lambda_j t} \mathbf{v}_j, \qquad (4.47)$$

where the $\{\mathbf{v}_j\}_{j=1}^{|\hat{\mathcal{X}}|}$ are the right eigenvectors of \mathbf{G} , and the $\{a_j\}_{j=1}^{|\hat{\mathcal{X}}|}$ are complex-valued coefficients derived from the initial distribution $\mathbf{p}(0)$. In this case, the solution's timevariation has an exponential character. This strongly suggests that, in our efforts to bound the solution, we should use basis functions which are also exponential with respect to time. Furthermore, it strongly suggests that the coefficients ρ appearing in these basis functions should match the eigenvalues of the matrix \mathbf{G} .

4.3.7 Choosing the Values of ρ

An obvious problem with the idea of choosing our values of ρ to match the eigenvalues of **G** is that there can be as many distinct eigenvalues as there are reachable states – often a huge number. Recall that each value of $\rho \in \mathcal{R}$ has an associated collection of decision variables $\tilde{\mathbf{z}}^{(\rho)}$ and constraints in SDP (4.41). It is not tractable to have such a large number of variables and constraints; so we can only hope to use some relatively small subset of the eigenvalues in defining the set \mathcal{R} .

This brings us to the question: which eigenvalues should we use? Our computational experience suggests that we should pick the values of ρ to approximate the real parts of the first several distinct eigenvalues of the matrix **G** when listed in order of increasing magnitude. By the construction of **G**, one of these eigenvalues is guaranteed to be zero, so we will always have $\rho = 0$ as one of our members of \mathcal{R} . Using Gershgorin's Circle Theorem[23], one can show that the nonzero eigenvalues of **G** all have strictly negative real parts.

The next question is: how can we approximate the eigenvalues we'd like to use in defining the set \mathcal{R} ? There are several possible strategies. First, it may be possible to construct a smaller instance of the problem we actually care about and use the eigenvalues of the smaller system as a proxy for the eigenvalues of the larger problem. For example, if the system we care about has an unmanageably large but finite number of reachable states, we could consider another instance of this system with the same reactions but with fewer molecules in the initial state \mathbf{x}_0 . Alternatively, it may be possible to decompose one large reaction system into two (smaller) independent

subsystems though the elimination of a single reaction. We can then approximate the eigenvalues of each smaller system separately, and then use their union to approximate the eigenvalues of the original system. The **G** matrix corresponding to one of these smaller systems might still be quite large. However, since it is also sparse, its low-magnitude eigenvalues can feasibly be approximated using an iterative Krylov subspace method[55]. This strategy is employed with measured success in Section 4.9. It may also be the case that the stoichiometry of the system is structured in such a way that it guarantees infinitely many reachable states, regardless of the initial state. In this case, the above strategy obviously won't work. An alternative strategy functions at these states, and setting \mathcal{R} to be the negatives of these values (being certain to include $\rho = 0$). This strategy is consistent with the observation that the relevant eigenvalues are usually the same order of magnitude as the reaction rate constants. However, it is clearly imprecise. Fortunately, as we will see in Section 4.5.3, the bounding method is tolerant of inaccuracy in the estimates of the eigenvalues.

4.3.8 Bounds on Higher-Order Moments

Tracing the argument leading to the formulation of SDP (4.41), we see that the exact identity of the objective function $\tilde{\mu}_{\mathbf{e}_i}(T)$ is not critical. This implies that we can replace $\tilde{\mu}_{\mathbf{e}_i}(T)$ with other moment-based expressions to obtain bounds on other quantities besides the mean molecular count of species *i*. For example, we can replace $\tilde{\mu}_{\mathbf{e}_i}(T)$ with any moment $\tilde{\mu}_{\mathbf{j}}(T)$ such that $|\mathbf{j}| \leq 2n$, to obtain bounds on $\mu_{\mathbf{j}}(T) = \langle \hat{\mathbf{X}}^{\mathbf{j}}(T) \rangle$. An example is given in Appendix C. More generally, we can readily bound any quantity that can be expressed as an affine function of the moments $\boldsymbol{\mu}(T)$, because the substitution of such a function into the objective function of SDP (4.41) gives an optimization problem that is still clearly an SDP.

4.3.9 Bounds on the Variance

In Section 4.2.6, we noted that the variance in the molecular count of species i at

time T can be expressed in terms of the moments $\boldsymbol{\mu}(T)$. In particular, $\sigma_i^2(T) = \mu_{2\mathbf{e}_i}(T) - \mu_{\mathbf{e}_i}^2(T)$. Given this situation, it is natural to wonder if one can obtain bounds on the variance of species *i* by simply substituting $\tilde{\mu}_{2\mathbf{e}_i}(T) - \tilde{\mu}_{\mathbf{e}_i}^2(T)$ in place of $\tilde{\mu}_{\mathbf{e}_i}(T)$ in SDP (4.41), giving Problem (4.48):

$$\sigma_{i}^{2}(T)^{U} = \max_{\substack{\boldsymbol{\tilde{\mu}}(T),\\ \boldsymbol{\tilde{z}}^{(\rho)}, \forall \rho \in \mathcal{R}}} \quad \tilde{\mu}_{2\mathbf{e}_{i}}(T) - \tilde{\mu}_{\mathbf{e}_{i}}^{2}(T)$$
s.t. $\tilde{\mu}_{0}(T) = 1,$
 $\tilde{\mu}(T) \in C_{n}(\boldsymbol{\alpha}, \boldsymbol{\beta}), \quad \langle \boldsymbol{\mu}(T) \in \boldsymbol{\mathcal{R}},$
 $\tilde{\mathbf{z}}^{(\rho)} \in C_{n}(\boldsymbol{\alpha}, \boldsymbol{\beta}), \quad \forall \rho \in \mathcal{R},$
Equation (4.31) holds, $\forall \rho \in \mathcal{R},$
Equation (4.33) holds, $\forall \rho \in \mathcal{R}.$

Since the expression $\tilde{\mu}_{2\mathbf{e}_i}(T) - \tilde{\mu}_{\mathbf{e}_i}^2(T)$ is not affine with respect to the moments $\boldsymbol{\mu}(T)$, Problem (4.48) is not an SDP, and it is not immediately obvious that we can solve it. However, it can be reformulated as an SDP in two steps:

- 1. Replace the objective with a "dummy variable", s, and add the constraint $s \leq \tilde{\mu}_{2\mathbf{e}_i}(T) \tilde{\mu}_{\mathbf{e}_i}^2(T)$.
- 2. Use the Schur complement lemma[71] to express this nonlinear inequality equivalently as the LMI

$$\begin{bmatrix} \tilde{\mu}_{2\mathbf{e}_i}(T) - s & \tilde{\mu}_{\mathbf{e}_i} \\ \tilde{\mu}_{\mathbf{e}_i} & 1 \end{bmatrix} \succeq \mathbf{0}.$$
 (4.49)

Doing so gives an SDP for calculating an upper bound on the variance in the molecular

count of species i at time T:

$$\sigma_{i}^{2}(T)^{U} = \max_{\substack{\tilde{\mu}(T), s, \\ \tilde{\mathbf{z}}^{(\rho)}, \forall \rho \in \mathcal{R}}} s$$
s.t.
$$\begin{bmatrix} \tilde{\mu}_{2\mathbf{e}_{i}}(T) - s & \tilde{\mu}_{\mathbf{e}_{i}} \\ \tilde{\mu}_{\mathbf{e}_{i}} & 1 \end{bmatrix} \succeq \mathbf{0},$$

$$\tilde{\mu}_{\mathbf{0}}(T) = 1,$$

$$\tilde{\mu}(T) \in C_{n}(\boldsymbol{\alpha}, \boldsymbol{\beta}),$$

$$\tilde{\mathbf{z}}^{(\rho)} \in C_{n}(\boldsymbol{\alpha}, \boldsymbol{\beta}), \quad \forall \rho \in \mathcal{R},$$
Equation (4.31) holds,
$$\forall \rho \in \mathcal{R},$$
Equation (4.33) holds,
$$\forall \rho \in \mathcal{R}.$$

It is important to note that changing the "max" to a "min" in SDP (4.50) does not result in a useful lower bound on variance. This is related to the fact that $\tilde{\mu}_{2\mathbf{e}_i}(T) - \tilde{\mu}_{\mathbf{e}_i}^2(T)$ is a concave function of $\tilde{\boldsymbol{\mu}}(T)$. Further details can be found in Chapter 3, in which we construct a similar SDP for bounding variances in the steadystate distribution of a stochastic chemical kinetic system.

4.3.10 Conservatism in the Bounds

As described in Chapter 3, there are several sources of conservatism in the bounds calculated by solving SDP (4.41) (and its variations). The first of these is related to the fact that our choice of m, the cut-off of what we consider to be a "low-order" moment, is somewhat arbitrary. The second source of conservatism is that the necessary conditions appearing in SDP (4.41) in no way reflect the physical constraint that the number of molecules of each species must be an integer. These sources of conservatism are discussed at length in Chapter 3, and the interested reader is referred there for further details.

There is, however, one source of conservatism which cannot be found in Chapter 3 and which is unique to the dynamic problem. This conservatism comes from our choice of the set \mathcal{R} . As we've already pointed out, Conditions (4.27) and (4.39)

hold for all $\rho \in \mathbb{R}$. However, for Problem (4.41) to be computationally tractable, we can only enforce these conditions for some finite subset $\mathcal{R} \subset \mathbb{R}$. In a sense, we are thus relaxing Conditions (4.27) and (4.39) for all $\rho \in \mathbb{R}$ such that $\rho \notin \mathcal{R}$. Doing so may introduce some conservatism in the resulting bounds. This suggests that adding elements $\rho \in \mathbb{R}$ to our set \mathcal{R} will improve the quality of the bounds. When we come to the examples in Section 4.5, we will see that this is, in fact, the case.

4.3.11 Scaling

As pointed out in Chapter 3, one shortcoming of moment-based SDPs such as Problem (4.41) is that they can give solvers numerical difficulties. This is especially true if the SDPs are not appropriately scaled. We discuss some strategies for scaling in Chapter 3, so we will not go into details here. However, we do wish to point out that, if one solves a sequence of bounding problems for increasing times T_j , the bounds at time T_j could be helpful in appropriately scaling the problem for time T_{j+1} .

4.3.12 Practical Applications

Sakurai and Hori[57] have demonstrated how moment-based SDPs for bounding the means and variances of steady-state distributions can be used in the design of biologically-inspired stochastic chemical kinetic systems called "biocircuits". Specifically, they use moment-based SDPs to synthesize a negative feedback biocircuit with statistical design specifications, and then conduct a sensitivity analysis on this system.

The dynamic bounding method we have described above can be used in much the same way. However, it allows for an even more sophisticated design analysis that also accounts for dynamics.

4.4 Toy Example

In this section, we apply SDPs (4.41) and (4.50) to a simple stochastic chemical kinetic systems as a proof of concept. Consider the simple irreversible reaction

$$A + B \xrightarrow{c_1} C \tag{4.51}$$

with rate constant $c_1 = 1$ s⁻¹, and known initial molecular counts of A = 3, B = 4, and C = 0. If we select A as the species to consider independent, this translates to an initial probability distribution $P(\cdot, 0) = \delta_3$, where all of the probability is concentrated on the reduced state $\hat{x} = 3$. This system has only four reduced states ($\hat{x} = 3, ..., 0$). Thus, we can compute the infinitesimal generator matrix for this system,

$$\mathbf{G} = \begin{bmatrix} -12 & 0 & 0 & 0\\ 12 & -6 & 0 & 0\\ 0 & 6 & -2 & 0\\ 0 & 0 & 2 & 0 \end{bmatrix},$$
(4.52)

and solve the CME directly. However, since this system features a bimolecular reaction, it exhibits the closure problem when subjected to a moment analysis. These features make this system a good first test for our bounding method.

4.4.1 Mean and Variance Bounds

If we repeatedly solve SDP (4.41) and its minimization counterpart for this system, taking m = 2 and $\mathcal{R} = \{0, -2\}$ (the first two eigenvalues of **G**), we obtain time-varying bounds on the mean molecular counts of each species. Similarly, if we repeatedly solve SDP (4.50) for this system, with the same \mathcal{R} and m, we can obtain time-varying upper bounds on variance for each molecular count. These bounds are shown in the top and bottom panels, respectively, of Figure 4-1. For comparison, we have also included the true means and variances computed by directly solving the CME. As expected, the mean bounds do, indeed, enclose the true means; and the variance upper bound does, indeed, exceed the true variance for all times t. This is consistent with the theory of Section (4.3).

4.4.2 Increasing the value of m

Equation (4.31), which appears as a constraint in SDPs (4.41) and (4.50), is derived from the expressions for the time derivatives of all moments up to order m. So, as we increase the value of m, the number of constraints implied by Equation (4.31) grows. Each of these additional constraints restricts the feasible set of SDPs (4.41) and (4.50). We might reasonably suspect, then, that increasing the value of m could lead to tighter bounds. This is, indeed, the case. If we recalculate the bounds shown in Figure 4-2 using m = 3 instead of m = 2, we obtain the slightly tighter bounds shown in Figure 4-2 (compare the plots at t = 4).

This is phenomenon is not specific to this example, but is a general feature of our bounding method. Increasing the value of the parameter m leads to monotonically tighter bounds (see Appendix C for a proof). But there is a trade-off, because increasing m also leads to larger, more numerically unstable SDPs. This same trade-off appeared in a slightly different form in the previous papers[13, 14, 56, 39, 24] on the moment-based steady-state bounding method. It is worth noting that the trade-off carries over to the dynamic extension. However, since the trade-off is well-documented in the steady-state papers and not unique to the dynamic bounding method, we discuss it no further here.

4.4.3 Using more values of ρ

In Section 4.3.10, we noted that the choice of the set \mathcal{R} can affect the quality of the resulting bounds. We demonstrate this by recalculating the bounds shown in Figure 4-2 with the enlarged set $\mathcal{R} = \{0, -2, -6\}$, including one more eigenvalue of **G**. The results are shown in Figure 4-3. The bounds are noticeably tighter for both the means and the variance, which is consistent with our prior reasoning.



Figure 4-1: Time-varying bounds on System (4.51), calculated using $\mathcal{R} = \{0, -2\}$ and m = 2. The points marked with circles and stars each correspond to the solution of an SDP and are theoretically guaranteed bounds. The lines interpolated between these points are not guaranteed bounds. They are included just to lead the eye. The dashed lines are the true values, obtained through direct solution of the CME. The bottom plot shows the variance in the molecular count of species A only, because the variances of the other species are identical.



Figure 4-2: Time-varying bounds on System (4.51), calculated using $\mathcal{R} = \{0, -2\}$ and m = 3. Note that the bounds are tighter than those shown in Figure 4-1. This is particularly true at long times.



Figure 4-3: This figure is equivalent to Figure 4-2 in every way, except that the bounds were calculated using the enlarged set $\mathcal{R} = \{0, -2, -6\}$, giving tighter bounds.



Figure 4-4: This figure is equivalent to Figure 4-2 in every way, except that the bounds were calculated using the enlarged set $\mathcal{R} = \{0, -2, -6, -12\}$, giving tighter bounds.

If we enlarge the set \mathcal{R} further by adding in $\rho = -12$, the final eigenvalue of \mathbf{G} , we obtain the bounds shown in Figure 4-4. The bounds on the mean molecular counts are not noticeably tighter, but the upper bound on the variance is.

In summary, the more values of the eigenspectrum of \mathbf{G} that we use in the set \mathcal{R} , the better the bounds become. However, even if we use the full eigenspectrum, the bounds are not perfect.

4.5 A Bit More Complexity

In this section, we apply SDPs (4.41) and (4.50) to a slightly more complex reaction system, where we've added a reversible reaction:

$$A + B \xrightarrow{c_1} C \xleftarrow{c_2} D \tag{4.53}$$

The rate constants for this system are $c_1 = 1 \text{ s}^{-1}$, $c_2 = 2.1 \text{ s}^{-1}$, and $c_3 = 0.3 \text{ s}^{-1}$ The initial molecular counts of A = 3, B = 4, C = 0, and D = 0. As with the previous example, this reaction system exhibits the closure problem when subjected to a moment analysis.

4.5.1 Mean and Variance Bounds

If we repeatedly solve SDP (4.41) and its minimization counterpart for this system, taking $\mathcal{R} = \{0, -2, -2.4\}$ and m = 3, we obtain time-varying bounds on the mean molecular counts of each species. Similarly, if we repeatedly solve SDP (4.50) for this system, with the same \mathcal{R} and m, we can obtain time-varying upper bounds on variance for each molecular count. These bounds are shown in the top and bottom panels, respectively, of Figure 4-5.

Again, this system is small enough that it is plausible to obtain the true mean and variance trajectories by numerically integrating the CME. These true trajectories are shown as dashed lines in Figure 4-5.

4.5.2 Using more values of ρ

As with the previous example, we now add a value of ρ to our set \mathcal{R} , repeat the bounding calculation, and see an improvement in the bounds. In particular, adding $\rho = -4.4$ to our \mathcal{R} , we obtain the bounds shown in Figure 4-6. Comparing with Figure 4-5, we see substantial improvement in the lower bound of the mean molecular count for species C. We also see that the limiting value of the variance upper bound for species C and D is about half of its previous value. Finally, for each species, the



Figure 4-5: Time-varying bounds on System (4.53), calculated using $\mathcal{R} = \{0, -2, -2.4\}$. The top plot shows bounds on mean molecular counts, while he bottom plot shows upper bounds on the variances in the counts. Species A is omitted because its behavior closely follows that of species B, and adding extra curves would only clutter the plot. The dashed lines are the true mean and variance trajectories, obtained by direct solution of the CME.


Figure 4-6: This figure is equivalent to Figure 4-5 in every way, except that the bounds were calculated using the enlarged set $\mathcal{R} = \{0, -2, -2.4, -4.4\}$, giving better results.

peak in the variance upper bound (around 0.5 s) has been reduced.

4.5.3 Sensitivity of the Values of ρ

As explained in Section 4.3.7, while any values of ρ will result in theoreticallyguaranteed bounds, we recommend picking the values of ρ to match the real parts of the first several distinct eigenvalues of the matrix **G**, when these eigenvalues are listed in order of increasing magnitude. This is exactly how we chose the values of ρ for the two foregoing examples. These two examples are small enough that we can calculate the eigenvalues directly. However, this will not be the case in general. Usually, the best we can hope for is some numerical approximation of the eigenvalues. This begs the question: how robust is our bounding method to the choice of ρ values? If the values of ρ are off by a little bit, do the bounds become so conservative that they are practically useless? To explore this idea, we repeated the bounding calculation for Reaction System (4.53), using a set of perturbed ρ values: $\mathcal{R} = \{0, -1.9, -2.6, -4.7\}$. The resulting bounds are shown in Figure 4-7. Comparing this plot with Figure 4-6, we see that the perturbation of the values of ρ did not substantially affect the quality of the computed bounds. We see that the perturbed ρ values create a slight long-time gap in the mean bounds for species C and D, which is undesirable. However, mean bounds on these species at intermediate times (e.g., t = 1 s) actually seem a little tighter. This demonstrates that the bounding method does not require exact knowledge of the eigenvalues of the underlying CME to obtain reasonable results.

That being said, the choice of ρ values does matter. Using a set of further perturbed ρ values ($\mathcal{R} = \{0, -6, -12, -18\}$), we produced the bounds shown in Figure 4-8. In this Figure, we see wide gaps in the long-time mean bounds for all species. Furthermore, the variance bounds are much less tight.

In summary, while the chosen values of ρ do not have to match low-magnitude eigenvalues **G** exactly, at least approximating them seems to be a good heuristic.

4.6 Complex Eigenvalues

Given our observation in Section 4.3.6 that it seems reasonable to choose the values of ρ to match the eigenvalues of the matrix **G**, it may seem odd that, in Section 4.3.7, we suggested focusing on only the real parts of these eigenvalues. In fact, if we know that some of the low-magnitude eigenvalues have nonzero imaginary parts, we can use this information to obtain tighter bounds.

For example, consider the cyclic system

$$A + B \xrightarrow{c_1} C$$

$$C \xrightarrow{c_2} D$$

$$D \xrightarrow{c_3} A + B$$

$$(4.54)$$

where the initial molecular counts are A = 2, B = 1, C = 1, and D = 0, and



Figure 4-7: This figure is equivalent to Figure 4-6, except that the bounds were calculated using the perturbed set $\mathcal{R} = \{0, -1.9, -2.6, -4.7\}$, giving slightly different results. In particular, notice the long-time gap that has appeared in the mean bounds for species C and D.



Figure 4-8: This figure is equivalent to Figures 4-6 and 4-7, except that the bounds were calculated using the further perturbed set $\mathcal{R} = \{0, -6, -12, -18\}$. This perturbation dramatically degrades the quality of the bounds.

the rate constants are $c_1 = 1 \text{ s}^{-1}$, $c_2 = 1.1 \text{ s}^{-1}$, and $c_3 = 0.9 \text{ s}^{-1}$. The smallestmagnitude eigenvalues of this system are $\lambda = 0, -2.1322 \pm 0.9741i, -4.1637 \pm 1.5837i$. If we follow the advice given in Section 4.3.7, and calculate bounds using $\mathcal{R} = \{0, -2.1322, -4.1637\}$, we obtain the bounds shown in the top panel of Figure 4-9. However, by making use of the knowledge of the imaginary parts of the eigenvalues, we can produce the slightly improved bounds shown in the bottom panel. The most notable improvements are for early times (t < 0.5 s)

Given this potential to improve the bounds by using the imaginary parts of the low-magnitude eigenvalues, why has this chapter been concerned almost solely with their real parts? The fact is "making use of the knowledge of the imaginary parts of the eigenvalues" is not trivial. One cannot simply use complex values of ρ in SDPs (4.41) and (4.50). The reason for this is that the argument for the derivation of LMIs (4.36) - (4.38) breaks down when ρ is complex-valued. It is possible to derive an analogous set of LMIs when ρ is complex-valued, but this requires introducing entirely new classes of decision variables and constraints. The resulting augmented versions of SDPs (4.41) and (4.50) are considerably more complicated. We felt that this extra complication would only distract from the main idea of this chapter, and, as demonstrated by Figure (4-9), it leads to only marginal improvement in the bounds. Accordingly, we have deferred the discussion of how to account for complex eigenvalues to Appendix C.



Figure 4-9: Bounds on the mean molecular counts of species A, C, and D of Reaction System (4.54). Bounds on species B are omitted, because they are similar to those shown for species A. The top panel shows bounds calculated without accounting for the imaginary components of the system's eigenvalues, while the bottom panel shows the slight improvement that can be achieved by accounting for these imaginary components. Both panels show the exact means calculated by directly solving the CME.

4.7 Perfect Bounds in the Absence of the Closure Problem

It is interesting to note that we can also apply our bounding method to stochastic chemical kinetic systems which do *not* exhibit the closure problem, and that, doing so, it is possible to obtain perfect bounds.

For example, consider the reaction system

$$A \xrightarrow{c_1} B \xrightarrow{c_2} C, \qquad (4.55)$$

where $c_1 = 1 \text{ s}^{-1}$, $c_2 = 3 \text{ s}^{-1}$, and there are initially 4 molecules of A and 0 molecules of each B and C. Since every reaction in this system is unimolecular, it does not exhibit the closure problem. The smallest-magnitude eigenvalues for this system are $\lambda = 0, -1, -3$. Solving SDP (4.41) and its minimization counterpart with $\mathcal{R} =$ $\{0, -1, -3\}$, we obtain the bounds shown in Figure 4-10. The upper and lower bounding curves are indistinguishable from one another because there is essentially no gap between them; they have collapsed upon the true mean trajectories.

This collapsing behavior is not unique to Reaction System (4.55). In Appendix C, we show that it can occur for any system which does not exhibit the closure problem, and whose matrix \mathbf{A}_L is diagonalizable (the typical case, in our experience). This feature supports the theoretical foundation of our bounding method. In particular, it suggests that our choice to use exponential basis functions is the appropriate choice.



Figure 4-10: Bounds on the mean molecular counts of species A, B, and C for Reaction System (4.55), which does not exhibit the closure problem. For this example, the bounding method calculates perfect bounds, collapsing on the true mean trajectories.

| State | $\mathbf{x} = (x_{\mathrm{A}}, x_{\mathrm{B}}, x_{\mathrm{C}}, x_{\mathrm{D}})$ |
|-------|---|
| 1 | (3,4,0,0) |
| 2 | (2, 3, 1, 0) |
| 3 | (2,3,0,1) |
| 4 | (1,2,2,0) |
| 5 | (1,2,1,1) |
| 6 | (1,2,0,2) |
| 7 | (0,1,3,0) |
| 8 | (0,1,2,1) |
| 9 | (0, 1, 1, 2) |
| 10 | (0,1,0,3) |

Table 4.1: The set of reachable states \mathcal{X} of the system described in Section 4.5.

4.8 Uncertainty in the Initial State

In each of the foregoing examples, we have assumed that we knew the initial molecular count exactly. This implies an initial probability distribution which is a Dirac distribution, where all the probability is concentrated on a single reachable state. However, as suggested in Section (4.3.2), our method can also handle the more general situation where we don't have exact knowledge of the initial molecular count, and the initial probability distribution (representing our knowledge of the system) is supported on several reachable states. We demonstrate this capability with the following example.

Again, consider Reaction System (4.53),

$$A + B \xrightarrow{c_1} C \xrightarrow{c_2} D,$$

with the same rate constants given in Section 4.5. In our prior analysis of this system, we assumed we knew the initial molecular counts A = 3, B = 4, C = 0, and D = 0. This implies the set of reachable states \mathcal{X} shown in Table 4.1. Furthermore, it implies an initial probability of zero for all states in Table 4.1, except State 1 which has an initial probability of one.

This time, we will assume uncertainty in the initial state, and we will express this uncertainty by assigning a nonzero initial probability to three distinct reachable states $\mathbf{x} \in \mathcal{X}$. In particular, we will assign initial probabilities of $\frac{1}{4}, \frac{1}{2}$, and $\frac{1}{4}$ to States 1, 4, and 10, respectively, with all other reachable states having an initial probability of zero. Once we have decided on the set of species to be considered independent (e.g., species A and C), we can easily calculate the initial low-order moments $\mu_L(0)$ corresponding to this initial distribution $P(\cdot, 0)$ using Equation (4.9). We can then apply SDPs (4.41) and (4.50) to calculate bounds on the means and variances for this system over time. For the sake of comparison to Figure 4-6, we again use m = 3 and $\mathcal{R} = \{0, -2, -2.4, -4.4\}$. The results are shown in Figure 4-11.



Figure 4-11: Bounds on the mean and variance for Reaction System (4.53) with an uncertain initial state.

The first thing to notice in comparing Figures 4-6 and 4-11 is that the starting point of each mean and variance trajectory is different between the two figures. This is consistent with the fact that the initial distribution $P(\cdot, 0)$ and thus the initial moments $\mathbf{y}_L(0)$ are different for the two figures. The second thing to notice is that both plots approach the same steady-state at long times. This is consistent with the fact that Reaction System (4.53) has just one steady-state, in which species C and D are in equilibrium. Finally, notice that the quality of the bounds is similar between the two plots. At least visually, the bounds in Figure 4-11 are just as tight as those in Figure 4-6. This may seem somewhat counter-intuitive given that Figure 4-11 was generated assuming uncertainty in the initial state. However, recall that once this uncertainty is expressed in an initial probability distribution $P(\cdot, 0)$, the means and variances (which are *expectation* values based on $P(\cdot, 0)$) are precisely defined.

4.9 A Larger Example

All of the examples given in the previous sections are small, in the sense that they contain very few reachable states. There are two reasons for this. First, the small size of these examples allows us to integrate numerically the CME and confirm that the true mean and variance trajectories do, in fact, lie within the bounds we have computed. Secondly, as suggested in Section 4.3.11, moment-based SDPs often lead to numerical difficulties, and these difficulties tend to scale with the size of the problem; using small examples allowed us to demonstrate various conceptual aspects of the bounding method without the distraction of numerical difficulties.

However, we do not wish the readers to be left with the impression that our bounding method is only applicable to small toy models. In theory, it could be applied to stochastic chemical kinetic systems of arbitrary size. Because the reachable states are not explicitly represented in our bounding SDPs, in theory, it does not matter whether the number of reachable states is ten, ten billion, or infinite.

To illustrate this point, we now apply our bounding method to a larger reaction system, taken from Chapter 3 and reproduced in Figure 4-12.

We assumed initial molecular counts of A = 53, F = 53, and zero for all other species. One can show[14] that this initial condition implies just over a billion reachable states (1 068 505 812, to be exact). Thus, solving the CME for this system is impractical. We can, however, apply our dynamic bounding method. If we repeatedly solve SDP (4.41) and its minimization counterpart for this system, taking $\mathcal{R} = \{0, -0.354, -0.651, -0.857, -1.062, -1.072, -1.374\}$ and m = 2, we obtain time-



Figure 4-12: A larger reaction system

varying bounds on the mean molecular counts of species A and H. These bounds are shown in Figure 4-13.

The reader may naturally wonder how we decided on the values to include in the set \mathcal{R} . After all, with one billion reachable states, computing the low-magnitude eigenvalues of the infinitesimal generator matrix **G** is a daunting proposition. What we did was consider a smaller instance of the reaction system, in which only 4 molecules of A and F were present initially. This system has only 570 reachable states, so estimating the low-magnitude eigenvalues of the corresponding **G** is manageable, via Arnoldi iteration[55]. These approximations were used in our set \mathcal{R} .

We have included this example to demonstrate that our bounding method can, in principle, be extended to very large systems. However, a few disclaimers are necessary. First, we are showing bounds for species A and H, because the bounds for these species are reasonably high-quality. However, the bounds for several other species in the reaction system are less impressive. Second, in solving the SDPs leading to the bounds in Figure 4-13, our solver of choice (SeDuMi[66]) often reported numerical problems. Accordingly, while our bounding method is useful for analyzing large systems in theory, we cannot yet claim that it is useful for analyzing such large systems in practice.



Figure 4-13: Bounds on the mean molecular counts of species A and H for the reaction system shown in Figure 4-12.

4.10 Comparison with SSA

To provide a comparison, Figure 4-13 also shows the averaged results of 100 runs of Gillespie's Stochastic Simulation Algorithm (SSA)[27], implemented in StochKit2[58]. While the sample mean trajectories apparently fall within our calculated bounds in this case, it is important to note that this result is not theoretically guaranteed; our bounds are for the true mean $\langle X(t) \rangle$, but the sample mean $\overline{X}(t)$ is a random variable, which can deviate substantially from the true mean. Still, since the sample mean converges to the true mean in the limit of large sample size, it is not surprising to see that it falls within our bounds.

The number of SSA runs was selected to provide a comparison of the capabilities of SSA and our bounding method, given equal CPU time. Specifically, the CPU time required to simulate the 100 SSA trajectories (130.9 s) is comparable to the CPU time required to solve the SDPs (125.9 s) which give the upper and lower bounds on the mean for species A at the final time point, t = 5 s. This equal-time comparison serves to highlight three differences between SSA and our bounding method. First, using SSA is inefficient for this reaction system because it contains rate constants of widely differing orders of magnitude (see Appendix B). On the other hand, while the stiffness of the system may affect the numerical conditioning of the bounding SDPs, it does not affect their theoretical solution time. Second, we revisit the point made in Section 4.3.5. To obtain information for this reaction system at t = T using SSA, we must simulate the entire time interval from t = 0 to t = T. On the other hand, with our bounding method, we must solve only the upper- and lower-bounding SDPs for t = T. We do not have to do any calculations concerned with intermediate time points. Finally, while our bounds are deterministic and will have the same values each time they are computed, the results of the SSA simulation will vary.

4.11 An Open System

For completeness, we also applied our bounding method to an open system with infinitely many reachable states:

where $c_1 = 1 \text{ s}^{-1}$, $c_2 = 0.01 \text{ s}^{-1}$, and the initial count of molecules is A = 2. To obtain values of ρ to use in the set \mathcal{R} , we followed the second strategy given in Section 4.3.7. Specifically we chose $x_A = 1, 2, 3, 4$ as our representative states, evaluated the propensity functions at these states, and then took the negatives of the resulting values, giving $\mathcal{R} = \{0, -0.02, -0.06, -0.12, -1\}$. If we repeatedly solve SDP (4.41) and its minimization counterpart for this system, using the \mathcal{R} specified above and m = 4, we obtain time-varying bounds on the mean shown in Figure 4-14. For comparison, we have also included the averaged results of 100 SSA runs.

4.12 Next Steps: Bounds on Probability

In Chapter 3, we also formulated SDPs for calculating an upper bound on the steady-state probability that the molecular count of species i is an arbitrary interval $[x_{\min}, x_{\max}]$, and we saw that this led to bounding histograms. We also noted that we



Figure 4-14: Bounds on the mean molecular counts of species A and H for the reaction system shown in Figure 4-12.

could bound the probability that the steady-state probability distribution assigns to an arbitrary basic semi-algebraic set, i.e., a set of the form

$$\{\hat{\mathbf{x}} \in \mathbb{R}^N : g_j(\hat{\mathbf{x}}) \ge 0, \ j = 1, \dots, K\},\tag{4.56}$$

where each $g_j(\cdot)$ for $j = 1, \ldots, K$ is a polynomial in $\mathbf{\hat{x}}$.

These ideas could readily be extended to the dynamic problem. For example, we could calculate bounds on the marginal distribution of species i at several times t. Such bounds have recently been shown to be useful in the analysis of individual cells[21]. Furthermore, we could bound the probability that this distribution assigns to an arbitrary basic semi-algebraic set over time. These extensions of the method are left for future work.

4.13 Conclusion

This chapter has described a method for calculating rigorous bounds on time-varying stochastic chemical kinetic systems. In particular, we have formulated SDPs for calculating time-varying bounds on the mean molecular count of each species in the system and the variances in these counts. This idea is an extension of the method described by several authors [13, 14, 56, 39, 24] for calculating bounds on the steadystate (i.e., stationary) distribution of a stochastic chemical kinetic system.

As a proof of concept, we have demonstrated the bounding method for a toy stochastic chemical system, with a single reaction. For this example, we have seen that our bounds are, in fact, valid. Furthermore, we have seen that they can be very tight, given the appropriate choice of the parameter set \mathcal{R} .

We also applied the bounding method to a slightly more complicated reaction system, which reaches a dynamic equilibrium in the limit of $t \to +\infty$. With this example, we saw that the bounds we obtain are not dramatically sensitive to the values of ρ we select in our parameter set \mathcal{R} .

While the majority of the chapter was written assuming that the parameter set \mathcal{R} contained strictly real values ρ , in Section 4.6 we saw that it is possible to obtain improved bounds by also using values of ρ with nonzero imaginary parts – though at the expense of solving a larger, more complicated SDP.

In Section 4.7, we saw an example which does not exhibit the closure problem, for which the bounds generated by our method collapse upon the true mean trajectories, supporting the theory underlying our approach.

In Section 4.8, we demonstrated that our method can also handle the scenario when the initial state of the system is not known exactly and we instead have nonzero initial probabilities associated with several reachable states.

In theory, our bounding method could be applied to stochastic chemical kinetic systems of arbitrary size. To demonstrate this point, in Section 4.9, we applied the method to a reaction system with more than a billion reachable states, though with qualified success. For our bounding method to be reliably useful in analyzing such large systems, there are two practical issues that must be overcome: first, we need to formalize a procedure for selecting the set \mathcal{R} ; second, we need to further explore options for mitigating the numerical issues mentioned in Sections 4.3.11 and 4.9. The first issue is left for future research. The second issue is the subject of Chapter 7.

Despite the method's incompleteness, it is a theoretically novel, interesting approach to the closure problem in stochastic chemical kinetics. We share it with the community in the hope that it might inspire further research in the area.

4.14 Implementation Details

All numerical examples in this chapter were computed on a 64-bit Dell Precision T3610 workstation with a 3.70 GHz Intel Xeon CPU. In the example, CVX [28] was used to model the SDP, using the default tolerance (i.e. "precision") settings. SeDuMi [66] was used as the underlying solver.

Chapter 5

Improved Bounds on Stochastic Chemical Kinetic Systems through an Alternative Representation of State Space

The contents of this chapter have not yet been published.

5.1 Introduction

As described in Chapter 3, one of our unique contributions to the bounding method was to show how it could be adapted to account for reaction invariants, which constrain the set of reachable states. However, after publishing several papers on the subject [13, 14, 16, 15] on the subject, we realized that we didn't take the idea far enough, and that our previously published methods provided unnecessarily loose bounds for specific types of reaction systems. In the present chapter, we identify the physical feature of reaction systems for which our previously published bounding method is ill-suited. We then describe an alternative formulation of the bounding method which addresses the problem. Finally, we demonstrate the advantages and limitations of the alternative formulation with numerical examples.

5.2 Notation

Much of the notation in this chapter is the same as in previous chapters. However, there are some subtle but important differences, specific to this chapter. For example, we now assign a physical interpretation to the auxiliary variables used in the definition the set of reachable states. Furthermore, we reserve the notation $P(\cdot, t)$ for the distribution over \mathcal{X} .

Let \mathbb{N} denote the natural numbers, $\{0, 1, 2, ...\}$. Consider a reacting chemical system with $N \in \mathbb{N}$ distinct molecular species and $R \in \mathbb{N}$ reactions, described by a stoichiometry matrix $\mathbf{S} \in \mathbb{Z}^{N \times R}$. Let the *state* of the system at time *t* be specified by the vector $\mathbf{X}(t) = (X_1(t), \ldots, X_N(t)) \in \mathbb{N}^N$, where $X_i(t) \ge 0$ is the count of molecules of species *i* present. Given an initial state $\mathbf{X}(0) \equiv \mathbf{x}_0$, the set of all *reachable states* is defined by

$$\mathcal{X} \equiv \left\{ \mathbf{x} \in \mathbb{N}^N : \begin{array}{c} \mathbf{x} = \mathbf{x}_0 + \mathbf{S}\boldsymbol{\epsilon}, \\ \mathbf{x} \ge \mathbf{0}, \quad \boldsymbol{\epsilon} \in \mathbb{N}^R \end{array} \right\},$$
(5.1)

where ϵ_r for $r \in \{1, \ldots, R\}$ is the *extent* of the *r*th reaction.

The probability associated with each reachable state $\mathbf{x} \in \mathcal{X}$ at time $t \geq 0$ is denoted $\Pr(\mathbf{X}(t) = \mathbf{x}) \equiv P(\mathbf{x}, t)$. The function $P(\cdot, t)$ is then the probability distribution over the set of reachable state \mathcal{X} at time t. In general, this probability distribution changes over time, and the way it changes is described by the *Chemical Master Equation* (CME):

$$\frac{dP}{dt}(\mathbf{x},t) = \sum_{r=1}^{R} \left[P(\mathbf{x} - \mathbf{s}_r, t) a_r(\mathbf{x} - \mathbf{s}_r) - P(\mathbf{x}, t) a_r(\mathbf{x}) \right], \quad \forall \mathbf{x} \in \mathcal{X}, \tag{5.2}$$

where \mathbf{s}_r is the *r*th column of the stoichiometry matrix, and $a_r(\cdot)$ is the *propensity* function associated with the *r*th reaction. When the reaction rates are governed by mass-action kinetics, these propensity functions are polynomials of their arguments. In particular, in the typical case, we have

$$a_{r}(\mathbf{x}) = \begin{cases} c_{r} \prod_{j=1}^{N} {x_{j} \choose \gamma_{j,r}}, & \text{if } \mathbf{x} \in \mathcal{X}, \\ 0, & \text{otherwise.} \end{cases}$$
(5.3)

where c_r is the rate constant associated with reaction r, and $\gamma_{j,r} = -\min\{s_{j,r}, 0\}$ is the (nonnegative) stoichiometric coefficient associated with species j in reaction r. See Gillespie [26] and Higham [31] for details.

Notice that $a_r(\mathbf{x})$ is defined such $a_r(\mathbf{x}) = 0$ if $x_j < \gamma_{j,r}$ (i.e., if the number of molecules of species j is less than the number required for the occurrence of reaction r). This prevents the flow of probability from a state $\mathbf{x} \in \mathcal{X}$ to a state $\mathbf{x}' \equiv \mathbf{x} + \mathbf{s}_r \in \mathbb{Z}^N$ which would be outside the set of reachable states, for the reason that $x'_j < 0$. In other words, for any state $\mathbf{x} \in \mathcal{X}$ and any reaction r such that $\mathbf{x} + \mathbf{s}_r \notin \mathcal{X}$, we have $a_r(\mathbf{x}) = 0$.

5.3 Previous Work

Again, most of what follows is review from previous chapters. This review will help us develop the novel material of this chapter by appealing to analogy with the previous material. Also, there are slight differences in notation. For example, we now reserve the symbol $\hat{P}(\cdot, t)$ for the distribution over $\hat{\mathcal{X}}$.

5.3.1 Invariants and Reduced Reachable States

In our previous work on bounding stochastic chemical kinetic systems [14, 16], we made use of the fact that reaction invariants can lead to a lower-dimensional representation of the state. In particular, the stoichiometry matrix \mathbf{S} often has a nontrivial

left null space, spanned by a set of basis vectors $\{\mathbf{b}_1, \ldots, \mathbf{b}_L\} \subset \mathbb{R}^N$. Setting

$$\mathbf{B} = \begin{bmatrix} \mathbf{b}_1^{\mathrm{T}} \\ \vdots \\ \mathbf{b}_L^{\mathrm{T}} \end{bmatrix}, \tag{5.4}$$

we noted that the the state $\mathbf{X}(t)$ must satisfy the affine equation

$$\mathbf{BX}(t) = \mathbf{Bx}_0 \equiv \mathbf{f} \tag{5.5}$$

for all times $t \ge 0$. We referred to the vector \mathbf{f} as the *invariant values*. Concatenating L linearly independent columns of \mathbf{B} into a matrix $\tilde{\mathbf{B}} \in \mathbb{R}^{L \times L}$, and collecting the remaining $N - L \equiv \hat{N}$ columns in a matrix $\hat{\mathbf{B}} \in \mathbb{R}^{L \times \hat{N}}$, we saw that the above equation could be written equivalently as

$$\tilde{\mathbf{B}}\tilde{\mathbf{X}}(t) + \hat{\mathbf{B}}\hat{\mathbf{X}}(t) = \mathbf{f}$$
(5.6)

where $\tilde{\mathbf{X}}(t) \in \mathbb{R}^{L}$ consists of the components of $\mathbf{X}(t)$ corresponding to the columns of $\tilde{\mathbf{B}}$, and $\hat{\mathbf{X}}(t) \in \mathbb{R}^{\hat{N}}$ is defined similarly. Rearranging (5.6), we obtained

$$\tilde{\mathbf{X}}(t) = \tilde{\mathbf{B}}^{-1}\mathbf{f} - \tilde{\mathbf{B}}^{-1}\hat{\mathbf{B}}\hat{\mathbf{X}}(t), \qquad (5.7)$$

which shows that specifying $\hat{\mathbf{X}}(t)$ determines the values of $\tilde{\mathbf{X}}(t)$. Accordingly, we labeled the species represented by $\hat{\mathbf{X}}(t)$ as the *independent species*, and those represented by $\tilde{\mathbf{X}}(t)$ as the *dependent species*. Since specifying $\hat{\mathbf{X}}(t)$ determines $\tilde{\mathbf{X}}(t)$, we saw that the state of the system can be expressed entirely in terms of the lowerdimensional $\hat{\mathbf{X}}(t)$. It follows that every full-dimensional reachable state $\mathbf{x} \in \mathcal{X} \subset \mathbb{N}^N$ can be expressed equivalently in terms of a *reduced reachable state* $\hat{\mathbf{x}} \in \mathbb{N}^{\hat{N}}$, obtained by extracting the components of the independent species from \mathbf{x} . We denoted the set of reduced reachable states using the symbol $\hat{\mathcal{X}} \subset \mathbb{N}^{\hat{N}}$.

5.3.2 A Reduced Probability Distribution

Having defined the set of reduced reachable states, $\hat{\mathcal{X}}$, we then defined a probability distribution over these reduced states, writing $\hat{P}(\hat{\mathbf{x}}, t) \equiv \Pr(\hat{\mathbf{X}}(t) = \hat{\mathbf{x}})$, for all $\hat{\mathbf{x}} \in \hat{\mathcal{X}}$. The way this probability distribution changes over time is given by a slight variation of the CME,

$$\frac{d\hat{P}}{dt}(\hat{\mathbf{x}},t) = \sum_{r=1}^{R} \left[\hat{P}(\hat{\mathbf{x}} - \hat{\mathbf{s}}_{r}, t) \hat{a}_{r}(\hat{\mathbf{x}} - \hat{\mathbf{s}}_{r}) - \hat{P}(\hat{\mathbf{x}}, t) \hat{a}_{r}(\hat{\mathbf{x}}) \right], \quad \forall \hat{\mathbf{x}} \in \hat{\mathcal{X}}, \tag{5.8}$$

in which $\hat{\mathbf{s}}_r$ is the reduced form of \mathbf{s}_r , containing only those components corresponding to the independent species; and $\hat{a}_r(\cdot)$ is a reduced form of $a_r(\cdot)$ in the sense that it accepts only the counts of the independent species as inputs. Assuming known invariant values \mathbf{f} and an initial probability distribution $\hat{P}(\cdot, 0)$, we defined the *steadystate* or *stationary* distribution of the system as $\hat{P}_{ss}(\cdot) \equiv \lim_{t \to +\infty} \hat{P}(\cdot, t)$.

5.3.3 Moments of the Reduced Distribution

We defined the moments of the distribution $\hat{P}(\cdot, t)$ using the equation

$$\mu_{\mathbf{j}}(t) \equiv \sum_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}} \hat{\mathbf{x}}^{\mathbf{j}} \hat{P}(\hat{\mathbf{x}}, t), \quad \forall \mathbf{j} \in \mathbb{N}^{\hat{N}}$$
(5.9)

where $\hat{\mathbf{x}}^{\mathbf{j}} \equiv \prod_{k=1}^{\hat{N}} \hat{x}_k^{j_k}$ is the monomial corresponding to the *multi-index* $\mathbf{j} \in \mathbb{N}^{\hat{N}}$. We used the symbol $\boldsymbol{\mu}(t)$ to denote the infinite sequence of moments at time t. The steady-state moments $\boldsymbol{\mu}_{ss}$ were defined similarly.

We noted that several physically meaningful statistical descriptions of $\hat{P}(\cdot, t)$ and $\hat{P}_{ss}(\cdot)$ could be expressed in terms their moments. For example, the mean molecular count of independent species $i \in \{1, \ldots, \hat{N}\}$ at time t can be written as $\langle \hat{X}_i(t) \rangle = \mu_{\mathbf{e}_i}(t)$, where \mathbf{e}_i is the *i*th coordinate vector.

5.3.4 The Bounding Paradigm

Next, we introduced our paradigm for obtaining rigorous bounds on statistical descriptions of the distribution $\hat{P}(\cdot, t)$ or $\hat{P}_{ss}(\cdot)$ which can be expressed in terms of their moments. For example, we suggested that we could calculate an upper bound on the mean molecular count of independent species $i \in \{1, \ldots, \hat{N}\}$ at steady state by solving an abstract optimization problem of the form

$$\langle \hat{X}_i \rangle_{\rm ss}^U \equiv \max_{\tilde{\mu}} \quad \tilde{\mu}_{\mathbf{e}_i}$$

s.t. $\tilde{\mu}$ satisfies necessary steady-state (5.10)
moment conditions.

Similarly, to calculate an upper bound on $\langle \hat{X}_i(t) \rangle$ for an arbitrary time $t \ge 0$, we had

$$\langle \hat{X}_{i}(t) \rangle^{U} \equiv \max_{\tilde{\boldsymbol{\mu}}(t)} \quad \tilde{\boldsymbol{\mu}}_{\mathbf{e}_{i}}(t)$$
s.t. $\tilde{\boldsymbol{\mu}}(t)$ satisfies necessary
moment conditions at time t .
$$(5.11)$$

We label the decision variable $\tilde{\mu}$ with a "tilde" to distinguish it from the vector of true moments μ defined in Equation (5.9). In principle, this vector of decision variables could be infinitely long, containing the full sequence of moments. However, in what follows, we will assume that the vector $\tilde{\mu}$ is finite and contains only as many moments as are required to express the necessary conditions that we will enforce as constraints.

5.3.5 Necessary Conditions on the Moment Sequences

The necessary conditions appearing in Problems (5.10) and (5.11) can be divided into three main categories, described in turn below.

Dynamics-Based Necessary Conditions

From the reduced form of the CME (5.8), one can derive a system of equations describing how the moments μ change over time. In particular, we can write

$$\frac{d\boldsymbol{\mu}_L}{dt}(t) = \mathbf{A}_L \boldsymbol{\mu}_L(t) + \mathbf{A}_H \boldsymbol{\mu}_H(t)$$
(5.12)

where μ_L is a vector containing the "low-order" moments up to some arbitrary order $m \in \mathbb{N}, \mu_H$ contains finitely many "higher-order" moments, and the constant matrices \mathbf{A}_L and \mathbf{A}_H are computed using the procedure described by Smadbeck and Kaznessis [63].

The dynamics-based necessary conditions are all derived from Equation (5.12). For example, for the steady-state analysis, we know that the time derivatives of the moments must be zero, leading to the equation

$$\mathbf{0} = \mathbf{A}_L \boldsymbol{\mu}_{L,\mathrm{ss}} + \mathbf{A}_H \boldsymbol{\mu}_{H,\mathrm{ss}}.$$
 (5.13)

For the transient analysis, we had linear constraints of the form

$$\boldsymbol{\mu}_{L}(t) - e^{\rho t} \boldsymbol{\mu}_{L}(0) = (\mathbf{A}_{L} - \rho \mathbf{I}) \mathbf{z}_{L}^{(\rho)} + \mathbf{A}_{H} \mathbf{z}_{H}^{(\rho)}$$
(5.14)

for each value of ρ in a finite set $\mathcal{R} \subset \mathbb{R}_{-}$, where $\mathbf{z}^{(\rho)}$ is a vector of generalized moments defined by

$$z_{\mathbf{j}}^{(\rho)} \equiv \int_0^t e^{\rho(t-\tau)} \mu_{\mathbf{j}}(\tau) d\tau, \quad \forall \mathbf{j} \in \mathbb{N}^{\hat{N}}, \forall \rho \in \mathcal{R}.$$
 (5.15)

The vectors $\mathbf{z}^{(\rho)}$ for each $\rho \in \mathcal{R}$ are also treated as decision variables in solving Problem (5.11).

Probability-Based Necessary Conditions

The first group of probability-based constraints came from the simple fact that the sum of all probabilities must always be one. For the steady-state analysis, this gave us

$$\mu_{0,ss} = 1.$$
 (5.16)

For the dynamic analysis, this gave us

$$\mu_{\mathbf{0}}(t) = 1 \tag{5.17}$$

and

$$z_{\mathbf{0}}^{(\rho)} = \begin{cases} T & \text{if } \rho = 0, \\ \frac{e^{\rho T} - 1}{\rho} & \text{otherwise,} \end{cases} \quad \forall \rho \in \mathcal{R}.$$
(5.18)

The second group of probability-based constraints was derived from the inequality

$$\sum_{\hat{\mathbf{x}}\in\mathcal{X}} q^2(\hat{\mathbf{x}}) \hat{P}(\hat{\mathbf{x}}, t) \ge 0, \quad \forall t \ge 0,$$
(5.19)

for any polynomial $q(\cdot)$. For the steady-state analysis, this lead to a *linear matrix* inequality (LMI)

$$\mathbf{M}_n^{\mathbf{0}}(\boldsymbol{\mu}_{\rm ss}) \succeq \mathbf{0},\tag{5.20}$$

where the matrix $\mathbf{M}_{n}^{\mathbf{0}}(\cdot)$, defined in Appendix B, is a linear function of its argument. Similarly, for the transient analysis, we had

$$\mathbf{M}_{n}^{\mathbf{0}}(\boldsymbol{\mu}(t)) \succeq \mathbf{0} \quad \text{and} \quad \mathbf{M}_{n}^{\mathbf{0}}(\mathbf{z}^{(\rho)}) \succeq \mathbf{0}, \ \forall \rho \in \mathcal{R}.$$
 (5.21)

Support-Based Necessary Conditions

Since the molecular counts of both the independent and dependent species must be nonnegative, it follows that the set of reduced reachable states $\hat{\mathcal{X}}$ must be contained in the polyhedral set

$$\bar{\mathcal{X}} \equiv \left\{ \hat{\mathbf{x}} \in \mathbb{R}^{\hat{N}} : \begin{array}{c} \hat{\mathbf{x}} \ge \mathbf{0} \\ \tilde{\mathbf{B}}^{-1} \mathbf{f} - \tilde{\mathbf{B}}^{-1} \hat{\mathbf{B}} \hat{\mathbf{x}} \ge \mathbf{0} \end{array} \right\},$$
(5.22)

Since $\hat{P}(\cdot, t)$ is nonzero only on $\hat{\mathcal{X}}$ and since $\hat{\mathcal{X}} \subset \bar{\mathcal{X}}$, it follows that the $\hat{P}(\cdot, t)$ is supported on $\bar{\mathcal{X}}$.

This led to a second group of LMIs based on the inequality

$$\sum_{\hat{\mathbf{x}}\in\mathcal{X}} g(\hat{\mathbf{x}})q^2(\hat{\mathbf{x}})\hat{P}(\hat{\mathbf{x}},t) \ge 0, \quad \forall t \ge 0,$$
(5.23)

where, again, $q(\cdot)$ is an arbitrary polynomial, and $g(\cdot)$ is one of the several polynomials defining $\bar{\mathcal{X}}$ through an inequality of the form $g(\hat{\mathbf{x}}) \geq 0$. For example, one such $g(\cdot)$ is $g(\hat{\mathbf{x}}) = \hat{x}_1$. We had two different types of LMIs derived from inequality (5.64), corresponding to the two different types of inequalities appearing in the definition of $\hat{\mathcal{X}}$. In particular, for the steady-state analysis, we had

$$\mathbf{M}_{n-1}^{\mathbf{e}_j}(\boldsymbol{\mu}_{\rm ss}) \succeq \mathbf{0}, \quad \forall j \in \{1, \dots, \hat{N}\},$$
(5.24)

and

$$\alpha_k \mathbf{M}_{n-1}^{\mathbf{0}}(\boldsymbol{\mu}_{\mathrm{ss}}) - \sum_{j=1}^{\hat{N}} \beta_{k,j} \mathbf{M}_{n-1}^{\mathbf{e}_j}(\boldsymbol{\mu}_{\mathrm{ss}}) \succeq \mathbf{0}, \quad \forall k \in \{1, \dots, L\}.$$
(5.25)

where α_k is the *k*th element of the vector $\tilde{\mathbf{B}}^{-1}\mathbf{f}$, and $\beta_{k,j}$ is the entry in the *k*th row and *j*th column of the matrix $\tilde{\mathbf{B}}^{-1}\hat{\mathbf{B}}$. For the transient analysis, we showed that the same LMIs must be satisfied by the moment vectors $\boldsymbol{\mu}(t)$ and $\mathbf{z}^{(\rho)}$ for each $\rho \in \mathcal{R}$.

5.3.6 Semidefinite Programs

Using all of the necessary conditions appearing in Section 5.3.5 as constraints in the abstract optimization problems of Section 5.3.4, we obtained two SDPs: one for the

steady-state problem,

$$\langle \hat{X}_i \rangle_{ss}^U \equiv \max_{\tilde{\boldsymbol{\mu}}} \quad \tilde{\mu}_{\mathbf{e}_i}$$
s.t. $\mathbf{A}_L \tilde{\boldsymbol{\mu}}_L + \mathbf{A}_H \tilde{\boldsymbol{\mu}}_H = 0,$
 $\tilde{\mu}_0 = 1,$
 $\mathbf{M}_n^0(\tilde{\boldsymbol{\mu}}) \succeq \mathbf{0},$
 $\mathbf{M}_{n-1}^{\mathbf{e}_j}(\tilde{\boldsymbol{\mu}}) \succeq \mathbf{0}, \quad \forall j \in \{1, \dots, \hat{N}\},$
 $\alpha_k \mathbf{M}_{n-1}^0(\tilde{\boldsymbol{\mu}}) - \sum_{j=1}^{\hat{N}} \beta_{k,j} \mathbf{M}_{n-1}^{\mathbf{e}_j}(\tilde{\boldsymbol{\mu}}) \succeq \mathbf{0}, \quad \forall k \in \{1, \dots, L\},$

$$(5.26)$$

and one for the transient problem,

$$\begin{split} \langle \hat{X}_{i}(t) \rangle^{U} &= \max_{\substack{\tilde{\mu}(t)\\ \tilde{\mathbf{z}}^{(\rho)}, \forall \rho \in \mathcal{R}}} \tilde{\mu}_{e_{i}}(t) \\ \text{s.t.} \quad \tilde{\mu}_{L}(t) - e^{\rho t} \tilde{\mu}_{L}(0) = (\mathbf{A}_{L} - \rho \mathbf{I}) \tilde{\mathbf{z}}_{L}^{(\rho)} + \mathbf{A}_{H} \tilde{\mathbf{z}}_{H}^{(\rho)}, \quad \forall \rho \in \mathcal{R}, \\ \tilde{\mu}_{0}(t) &= 1, \\ \tilde{z}_{0}^{(\rho)} &= \begin{cases} T & \text{if } \rho = 0, \\ \frac{e^{\rho T - 1}}{\rho} & \text{otherwise,} \end{cases} \quad \forall \rho \in \mathcal{R}, \\ \mathbf{M}_{n}^{0}(\tilde{\mathbf{\mu}}(t)) \succeq \mathbf{0}, \\ \mathbf{M}_{n}^{0}(\tilde{\mathbf{z}}^{(\rho)}) \succeq \mathbf{0}, \quad \forall \rho \in \mathcal{R}, \\ \mathbf{M}_{n-1}^{e_{j-1}}(\tilde{\mu}(t)) \succeq \mathbf{0}, \quad \forall j \in \{1, \dots, \hat{N}\}, \\ \mathbf{M}_{n-1}^{e_{j-1}}(\tilde{\mathbf{z}}^{(\rho)}) \succeq \mathbf{0}, \quad \forall j \in \{1, \dots, \hat{N}\}, \quad \forall \rho \in \mathcal{R}, \\ \alpha_{k}\mathbf{M}_{n-1}^{0}(\tilde{\mathbf{z}}^{(\rho)}) - \sum_{j=1}^{\hat{N}} \beta_{k,j}\mathbf{M}_{n-1}^{e_{j-1}}(\tilde{\mathbf{\mu}}(t)) \succeq \mathbf{0}, \quad \forall k \in \{1, \dots, L\}, \\ \alpha_{k}\mathbf{M}_{n-1}^{0}(\tilde{\mathbf{z}}^{(\rho)}) - \sum_{j=1}^{\hat{N}} \beta_{k,j}\mathbf{M}_{n-1}^{e_{j-1}}(\tilde{\mathbf{z}}^{(\rho)}) \succeq \mathbf{0}, \quad \forall k \in \{1, \dots, L\}, \\ \forall \rho \in \mathcal{R}. \end{cases}$$

$$(5.27)$$

Parameters

For the steady-state problem, there are two parameters. The first of these is $m \in \mathbb{N}$, our arbitrary cut-off of what we consider a "low-order" moment. The choice of maffects the length of the $\tilde{\mu}_L$ and $\tilde{\mu}_H$ vectors. Accordingly, it affects the number of rows and columns of the \mathbf{A}_L and \mathbf{A}_H matrices. The second parameter is $n \in \mathbb{N}$, which determines the size of the $\mathbf{M}_n^0(\tilde{\mu})$, $\mathbf{M}_{n-1}^0(\tilde{\mu})$, and $\mathbf{M}_{n-1}^{\mathbf{e}_j}(\tilde{\mu})$ matrices, as described in Appendix B. In principle, any combination of these two parameters gives an SDP which will yield valid bounds. However, for reasons explained in Chapter 3, we recommend linking the two parameter values through the equation $n = \lceil (m + q - 1)/2 \rceil$, where $q \in \mathbb{N}$ is the maximum reaction order in the underlying chemical system (e.g., for a system containing at most bimolecular reactions, q = 2). With the equation $n = \lceil (m + q - 1)/2 \rceil$, we can treat m as the sole parameter for the SDP. As explained in Chapters 3 and 4, the quality of the bounds increases monotonically with the values of m. However, larger values of m also yield larger SDPs, which are more prone to numerical instability. Thus, there is a trade-off in selecting m.

The transient problem also has $m \in \mathbb{N}$ and $n \in \mathbb{N}$ as parameters, and everything we've said in the previous paragraph applies. Furthermore, the choice of m affects the size of the $\tilde{\mathbf{z}}_{L}^{(\rho)}$ and $\tilde{\mathbf{z}}_{H}^{(\rho)}$ for each value of $\rho \in \mathcal{R}$. This set $\mathcal{R} \subset \mathbb{R}_{-}$ is also treated as a parameter. Any choice of \mathcal{R} leads to valid bounds. However, as we showed in Chapter 4, the quality of the bounds tends to be better when the values of \mathcal{R} approximate the most significant time constants of the underlying chemical system. In that same chapter, we gave some heuristics for picking \mathcal{R} . One of these heuristics was to pick several representative states $\hat{\mathbf{x}} \in \hat{\mathcal{X}}$ and then evaluate the propensity functions for these states. Adding values to the set \mathcal{R} can only improve the quality of the resulting bounds. However, again, there is a trade-off, because the size of SDP (5.27) increases with the cardinality of \mathcal{R} .

Bounds on other Statistical Descriptions

Solving SDPs (5.26) and (5.27) gives an upper bound on the mean molecular count of independent species *i* at steady state and at an arbitrary time $t \ge 0$, respectively. However, as we showed Chapters 3 and 4, variations on this theme can be used to bound other statistical descriptions of the unknown probability distributions $\hat{P}_{ss}(\cdot)$ and $\hat{P}(\cdot, t)$.

5.3.7 $\bar{\mathcal{X}}$ as an Outer-Approximation of $\hat{\mathcal{X}}$

As described in Chapter 3, the polyhedral superset $\bar{\mathcal{X}}$ is an outer-approximation of the set of reduced reachable states $\hat{\mathcal{X}}$ in the sense that $\bar{\mathcal{X}}$ does not require molecular counts to be integers (i.e., a point $\hat{\mathbf{x}} \in \bar{\mathcal{X}}$ doesn't necessarily satisfy $\hat{\mathbf{x}} \in \mathbb{N}^{\hat{N}}$ or $\tilde{\mathbf{B}}^{-1}\mathbf{f} - \tilde{\mathbf{B}}^{-1}\hat{\mathbf{B}}\hat{\mathbf{x}} \in \mathbb{N}^{L}$ while a point $\hat{\mathbf{x}} \in \hat{\mathcal{X}}$ does). We noted that this outer-approximation has potential to introduce conservatism into the bounds produced by solving the aforementioned SDPs. While this conservatism was slight, it is still undesirable. We described a mitigation strategy involving selectively adding LMIs to the bounding SDPs.

5.4 A Second Type of Outer-Approximation

In fact, there is a second way in which the set $\tilde{\mathcal{X}}$ may be an outer-approximation of $\hat{\mathcal{X}}$, leading to more dramatic conservatism in the bounds.

5.4.1 A Pathological Example

This second type of outer-approximation is best illustrated through an example. Consider the reaction system

$$\begin{array}{c} A + B \xrightarrow{c_1} C \\ A \xrightarrow{c_2} D \end{array} \tag{5.28}$$

with $c_1 = c_2 = 1 \text{ s}^{-1}$. The stoichiometry matrix for this system is

$$\mathbf{S} = \begin{bmatrix} -1 & -1 \\ -1 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}.$$
 (5.29)

The left null space of \mathbf{S} is spanned by the rows of

$$\mathbf{B} = \begin{bmatrix} 1 & 0 & 1 & 1 \\ 0 & 1 & 1 & 0 \end{bmatrix}.$$
 (5.30)

It follows that the sums $x_{\rm A} + x_{\rm C} + x_{\rm D}$ and $x_{\rm B} + x_{\rm C}$ are invariants of this reaction system. The values of these invariants depend on the initial condition $\mathbf{x}(0)$. Suppose that the initial condition is $\mathbf{x}(0) = (2, 3, 3, 0)$. Then, the values of the invariants are $\mathbf{B}^{\rm T}\mathbf{x}(0) = (5, 6)$. Selecting species A and B as the independent species, Equation (5.22) simplifies to

$$\bar{\mathcal{X}} = \left\{ \begin{bmatrix} x_A \\ x_B \end{bmatrix} \in \mathbb{R}^{\hat{N}} : \begin{bmatrix} x_A \\ x_B \end{bmatrix} \ge \mathbf{0} \\ \begin{bmatrix} 6 - x_B \\ -1 - x_A + x_B \end{bmatrix} \ge \mathbf{0} \right\}.$$
(5.31)

How does this $\bar{\mathcal{X}}$ compare with $\hat{\mathcal{X}}$? Inspecting Reaction System (5.28) and recalling the initial state $\mathbf{x}(0)$, we see that the set of full-dimensional states is

$$\mathcal{X} = \left\{ \begin{bmatrix} 2\\3\\3\\0 \end{bmatrix}, \begin{bmatrix} 1\\2\\4\\0 \end{bmatrix}, \begin{bmatrix} 1\\3\\3\\1 \end{bmatrix}, \begin{bmatrix} 0\\2\\4\\1 \end{bmatrix}, \begin{bmatrix} 0\\3\\3\\2 \end{bmatrix}, \begin{bmatrix} 0\\1\\5\\0 \end{bmatrix} \right\},$$
(5.32)

from which it follows trivially that the set of reduced reachable states is

$$\hat{\mathcal{X}} = \left\{ \begin{bmatrix} 2\\3 \end{bmatrix}, \begin{bmatrix} 1\\2 \end{bmatrix}, \begin{bmatrix} 1\\3 \end{bmatrix}, \begin{bmatrix} 0\\2 \end{bmatrix} \begin{bmatrix} 0\\3 \end{bmatrix} \begin{bmatrix} 0\\1 \end{bmatrix} \right\},$$
(5.33)

We can easily verify that $\hat{\mathcal{X}} \subset \overline{\mathcal{X}}$. However, notice that there are many integer-valued points $\hat{\mathbf{x}} \in \overline{\mathcal{X}}$ that are not elements of $\hat{\mathcal{X}}$. For example, each of the points in the set

$$S \equiv \left\{ \begin{bmatrix} 1\\ 4 \end{bmatrix}, \begin{bmatrix} 0\\ 6 \end{bmatrix}, \begin{bmatrix} 4\\ 5 \end{bmatrix} \right\}$$
(5.34)

is an element of $\bar{\mathcal{X}}$ but not $\hat{\mathcal{X}}$. Inspecting Reaction System (5.28) and the initial state $\mathbf{x}(0) = (2, 3, 3, 0)$, it is obvious that each of the states $\hat{\mathbf{x}} \in S$ cannot be in $\hat{\mathcal{X}}$; achieving these states would require the generation of A or B molecules, for which there is no mechanism. Yet, clearly, $S \subset \bar{\mathcal{X}}$. This leads us to the conclusion that, in this case, $\bar{\mathcal{X}}$ is a gross outer-approximation of $\hat{\mathcal{X}}$.

5.4.2 Implications for the Bounds

As we might expect based on the discussion in Section 5.3.7, this outer-approximation leads to undesirable conservatism in the bounds produced by solving the SDPs described in Section 5.3.6.

Steady-State Analysis

For example, if we try to upper bound the steady-state mean molecular count of species B by solving SDP (5.26), with any $m \in \mathbb{N}$, we obtain $\langle B \rangle_{ss}^U = 6$. This upper bound is valid, but it is also obviously conservative, given that the stoichiometry and initial condition of the reaction system require that $x_B(t) \leq 3$ for all $t \geq 0$.

The explanation for this conservatism in the upper bound has two parts. First, the LMIs appearing as constraints in SDP (5.26) are necessary conditions for a moment sequence $\boldsymbol{\mu}$ of a distribution supported on $\bar{\mathcal{X}}$. As we saw in Section 5.4.1, while the point $\hat{\mathbf{x}}^* \equiv (0, 6)$ is not an element of $\hat{\mathcal{X}}$, it is an element of $\bar{\mathcal{X}}$. Accordingly, the

Dirac distribution $\delta_{\hat{\mathbf{x}}^*}$, in which all of the probability is concentrated on the point $\hat{\mathbf{x}}^*$, is supported on $\bar{\mathcal{X}}$. It follows that the moments $\boldsymbol{\mu}$ of this distribution satisfy the LMIs of SDP (5.26). They also trivially satisfy the constraint $\mu_0 = 1$. Second, the point $\hat{\mathbf{x}}^* = (0, 6)$ implies 0 molecules of A. Inspecting Reaction System (5.28), we see that having 0 molecules of A implies that no reaction is possible (the propensity functions evaluate to zero). It follows that the distribution $\delta_{\hat{\mathbf{x}}^*}$, when subjected to the CME for this system, is invariant with time. Furthermore, its moments $\boldsymbol{\mu}$ do not change with time. In other words, $\boldsymbol{\mu}$ satisfies the first constraint of SDP (5.26). We've come to the conclusion, then, that $\boldsymbol{\mu}$ is a feasible point of SDP (5.26). This implies that the optimal value of SDP (5.26) is greater than or equal to the objective value corresponding to $\boldsymbol{\mu}$:

$$\mu_{\mathbf{e}_2} = \sum_{\hat{\mathbf{x}}\in\bar{\mathcal{X}}} \hat{\mathbf{x}}^{\mathbf{e}_2} \delta_{\hat{\mathbf{x}}^*} = \sum_{\hat{\mathbf{x}}\in\bar{\mathcal{X}}} \hat{x}_2 \delta_{\hat{\mathbf{x}}^*} = \hat{x}_2^* = 6$$
(5.35)

Thus, the inclusion of the point $\hat{\mathbf{x}}^*$ in the set $\bar{\mathcal{X}}$ artificially pushes the upper bound upward, degrading the quality of the bound.

Transient Analysis

What about the transient analysis? To solve SDP (5.27), we need to choose a set \mathcal{R} . We will do so using the heuristic described in Section 5.3.6, picking (somewhat arbitrarily) $\hat{\mathbf{x}} = (2,3)$ and $\hat{\mathbf{x}} = (1,2)$ as the "representative" reachable states. Evaluating the and summing the propensity functions for each of these states suggests using ρ values of -8 and -3. Including the standard value of $\rho = 0$, leaves us with $\mathcal{R} = \{0, -3, -8\}$.

Solving SDP (5.27) with this \mathcal{R} and m = 2 for a range of t values from t = 0 s to t = 3 s, we obtain the bounds shown in Figure 5-1. This system is small enough that directly solving the CME is also practical. Accordingly the exact CME solution is also plotted in Figure 5-1 for comparison.



Figure 5-1: Time-varying bounds on the mean molecular counts of species A, B, and D for Reaction System (5.28). These bounds were calculated by solving SDP (5.27) (and its minimization counterpart) over a range of t values, using m = 2 and $\mathcal{R} = \{0, -3, -8\}.$

Inspecting Figure 5-1, we see that the computed bounds are quite tight. At least, they are not more conservative than the bounds implied by stoichiometry, as we saw for the bound in the steady-state analysis. However, given that many of the LMIs appearing as constraints in SDP (5.27) are derived from $\bar{\mathcal{X}}$, we might wonder if the over-approximation built into $\bar{\mathcal{X}}$ is causing unnecessary conservatism in the bounds. In what follows, we will see that this is, in fact, the case.

5.4.3 Generalization

What went wrong in the above example? We've seen that the immediate problem is that the set $\bar{\mathcal{X}}$ is a substantial over-approximation of the true reachable set $\hat{\mathcal{X}}$, in the sense that there exist many integer-valued points $\hat{\mathbf{x}} \in \mathbb{N}^{\hat{\mathcal{N}}}$ such that $\hat{\mathbf{x}} \in \bar{\mathcal{X}}$ but $\hat{\mathbf{x}} \notin \hat{\mathcal{X}}$. However, the fundamental problem is that the set $\bar{\mathcal{X}}$ described in Equation (5.31) does not reflect the fact that the reactions appearing in Reaction System (5.28) are irreversible.

As the reader may guess, the problem is not specific to this example. The general expression for $\bar{\mathcal{X}}$ given in Equation (5.22) implicitly assumes reversible reactions.

Thus, if irreversible reactions are present, and this leads to $\bar{\mathcal{X}}$ being a substantial over-approximation of the true reachable set $\hat{\mathcal{X}}$, applying the bounding methods of Section 5.3 may lead to the same type of undesirable conservatism we've seen in the present pathological example.

5.5 An Alternative State Representation

Fortunately, there is a relatively simple way to fix the problem identified in the previous section. The key idea is to express the state of the system not in terms of the molecular counts $\mathbf{X}(t) \in \mathbb{N}^N$ but in terms of the extents of reaction $\mathbf{E}(t) \in \mathbb{N}^R$. Then, in analogy with our definition of \mathcal{X} , we can redefine the set of reachable states as

$$\mathcal{E} \equiv \left\{ \boldsymbol{\epsilon} \in \mathbb{N}^R : \mathbf{x}_0 + \mathbf{S}\boldsymbol{\epsilon} \ge \mathbf{0} \right\}.$$
(5.36)

Now, as before, we'd like to construct a polyhedral superset guaranteed to contain the set of reachable states. This superset is given by

$$\bar{\mathcal{E}} \equiv \left\{ \boldsymbol{\epsilon} \in \mathbb{R}^R : \begin{array}{c} \mathbf{x}_0 + \mathbf{S}\boldsymbol{\epsilon} \ge \mathbf{0}, \\ \boldsymbol{\epsilon} \ge \mathbf{0} \end{array} \right\}.$$
(5.37)

The first inequality appearing in this definition of $\overline{\mathcal{E}}$ reflects the physical fact molecular counts must be nonnegative, and is analogous to the two inequalities appearing in our definition of $\overline{\mathcal{X}}$. However, the second equality, $\epsilon \geq 0$, has no analogue in the definition of $\overline{\mathcal{X}}$; it expresses precisely what the definition of $\overline{\mathcal{X}}$ lacks – namely, that each elementary reaction is irreversible.

5.5.1 Redundancy in the State Representation

One undesirable feature of this state representation is that it may be redundant, in the sense that there may be multiple states $\boldsymbol{\epsilon} \in \mathcal{E}$ corresponding to the same state $\mathbf{x} \in \mathcal{X}$. This will be true whenever there exists some $\boldsymbol{\epsilon}_{\perp} \in \mathbb{N}^{R}$ such that $\boldsymbol{\epsilon}_{\perp} \neq \mathbf{0}$ and $\mathbf{S}\boldsymbol{\epsilon}_{\perp} = \mathbf{0}$. Then, for every $\boldsymbol{\epsilon} \in \mathcal{E}$ and the corresponding state $\mathbf{x} = \mathbf{x}_{0} + \mathbf{S}\boldsymbol{\epsilon} \in \mathcal{X}$, there is an infinite set of states

$$\{\boldsymbol{\epsilon} + \alpha \boldsymbol{\epsilon}_{\perp} : \alpha \in \mathbb{Z}, \ \boldsymbol{\epsilon} + \alpha \boldsymbol{\epsilon}_{\perp} \ge \mathbf{0}\} \subset \mathcal{E}$$

$$(5.38)$$

corresponding to \mathbf{x} .

For example, consider the simple reversible reaction system $A \Longrightarrow B$. The stoichiometry matrix for this system is

$$\mathbf{S} = \begin{bmatrix} -1 & 1\\ 1 & -1 \end{bmatrix}. \tag{5.39}$$

The vector $\boldsymbol{\epsilon}_{\perp} = (1, 1)$ satisfies $\mathbf{S}\boldsymbol{\epsilon}_{\perp} = \mathbf{0}$. In words, this means that the two reactions "cancel each other out". If the system is in state $\mathbf{x} \in \mathcal{X}$, and both reactions occur in sequence, incrementing the values of ϵ_1 and ϵ_2 , the system returns to the same \mathbf{x} , but has arrived at a new $\boldsymbol{\epsilon} \in \mathcal{E}$. Accordingly, for any $\mathbf{x} \in \mathcal{X}$ in this reaction system, there are infinitely many corresponding states $\boldsymbol{\epsilon} \in \mathcal{E}$.

In general, we will define the set of states $\epsilon \in \mathcal{E}$ corresponding to a state $\mathbf{x} \in \mathcal{X}$ as

$$\mathcal{E}(\mathbf{x}) \equiv \{ \boldsymbol{\epsilon} \in \mathcal{E} : \mathbf{x} = \mathbf{x}_0 + \mathbf{S}\boldsymbol{\epsilon} \}, \quad \forall \mathbf{x} \in \mathcal{X}.$$
(5.40)

Similarly, we can define $\mathcal{E}(\hat{\mathbf{x}})$ as the set of states $\boldsymbol{\epsilon} \in \mathcal{E}$ corresponding to a reduced state $\hat{\mathbf{x}} \in \hat{\mathcal{X}}$. From Equations (5.1) and (5.36) it follows that

$$\mathcal{E} = \bigcup_{\mathbf{x} \in \mathcal{X}} \mathcal{E}(\mathbf{x}).$$
(5.41)

Since there is a one-to-one correspondence between the elements of \mathcal{X} and $\hat{\mathcal{X}}$, it is also true that

$$\mathcal{E} = \bigcup_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}} \mathcal{E}(\hat{\mathbf{x}}).$$
(5.42)

The redundancy we have described in this section will require us to exercise some care in the derivations that follow, but it does not invalidate the representation of the state of the system in terms of extents.

5.6 Alternative Formulation of the Bounding Method

Having redefined the set of reachable states in terms of extents of reaction, \mathcal{E} , and having constructed the corresponding polyhedral superset $\overline{\mathcal{E}}$, we have apparently overcome the fundamental problem described in Section 5.4. In this section, we describe how the bounding method described in our previous work [14, 16] can be adapted to make use of this alternative representation of the state.

5.6.1 A New Probability Distribution

First, we can define a probability distribution over \mathcal{E} in much the same way that we defined the probability distribution over \mathcal{X} and $\hat{\mathcal{X}}$:

$$P_{\mathcal{E}}(\boldsymbol{\epsilon}, t) \equiv \Pr(\mathbf{E}(t) = \boldsymbol{\epsilon}), \quad \forall \boldsymbol{\epsilon} \in \mathcal{E}.$$
(5.43)

The function can be extended to the domain \mathbb{Z}^R if we specify that $P_{\mathcal{E}}(\boldsymbol{\epsilon}, t) \equiv 0$ for all $\boldsymbol{\epsilon} \in \mathbb{Z}^R$ such that $\boldsymbol{\epsilon} \notin \mathcal{E}$.

The distribution $P_{\mathcal{E}}(\cdot, t)$ is related to $P(\cdot, t)$ and $\hat{P}(\cdot, t)$ through the following equations:

$$P(\mathbf{x},t) = \sum_{\boldsymbol{\epsilon} \in \mathcal{E}(\mathbf{x})} P_{\mathcal{E}}(\boldsymbol{\epsilon},t), \quad \forall \mathbf{x} \in \mathcal{X},$$
(5.44)

$$\hat{P}(\hat{\mathbf{x}},t) = \sum_{\boldsymbol{\epsilon} \in \mathcal{E}(\hat{\mathbf{x}})} P_{\mathcal{E}}(\boldsymbol{\epsilon},t), \quad \forall \hat{\mathbf{x}} \in \hat{\mathcal{X}}.$$
(5.45)

5.6.2 Moments

Second, we can define moments of the distribution $P_{\mathcal{E}}(\cdot, t)$, in analogy with how we defined the moments of the distribution $\hat{P}(\cdot, t)$:

$$y_{\mathbf{j}}(t) \equiv \sum_{\boldsymbol{\epsilon} \in \mathcal{E}} \boldsymbol{\epsilon}^{\mathbf{j}} P_{\mathcal{E}}(\boldsymbol{\epsilon}, t), \quad \forall \mathbf{j} \in \mathbb{N}^{R}.$$
 (5.46)

As before, the zeroth order moment is nothing more than the sum of probabilities across all $\epsilon \in \mathcal{E}$. Accordingly, we have $y_0(t) = 1$, for all $t \ge 0$.

The moments $\mathbf{y}(t)$ can be used to express the same statistical descriptions of the reaction system that could be expressed in terms of the moments $\boldsymbol{\mu}(t)$. For example, the mean molecular count of species $i \in \{1, \ldots, N\}$ at time t is equal to

$$\langle X_i(t) \rangle = \sum_{\boldsymbol{\epsilon} \in \mathcal{E}} \left(x_{0,i} + \sum_{r=1}^R s_{i,r} \boldsymbol{\epsilon}_r \right) P_{\mathcal{E}}(\boldsymbol{\epsilon}, t),$$

$$= x_{0,i} \sum_{\boldsymbol{\epsilon} \in \mathcal{E}} P_{\mathcal{E}}(\boldsymbol{\epsilon}, t) + \sum_{r=1}^R s_{i,r} \sum_{\boldsymbol{\epsilon} \in \mathcal{E}} \boldsymbol{\epsilon}_r P_{\mathcal{E}}(\boldsymbol{\epsilon}, t),$$

$$= x_{0,i} y_{\mathbf{0}}(t) + \sum_{r=1}^R s_{i,r} y_{\mathbf{e}_r}(t),$$

$$= x_{0,i} + \sum_{r=1}^R s_{i,r} y_{\mathbf{e}_r}(t).$$

$$(5.47)$$

Moreover, any moment $\mu_{\mathbf{j}}(t)$, $\mathbf{j} \in \mathbb{N}^{\hat{N}}$ can be expressed as a linear combination of the moments $\mathbf{y}(t)$. We leave the derivation to Appendix D.

5.6.3 The Bounding Paradigm

The new bounding paradigm is essentially the same that we saw in Section 5.3.4. For example, to calculate an upper bound on the mean molecular count of species $i \in \{1, ..., N\}$ at steady-state and at an arbitrary time t, we solve the abstract optimization problems

$$\langle X_i \rangle_{\rm ss}^U \equiv \max_{\tilde{\mathbf{y}}} \quad x_{0,i} + \sum_{r=1}^R s_{i,r} \tilde{y}_{\mathbf{e}_r}$$
s.t. $\tilde{\mathbf{y}}$ satisfies necessary steady-state moment conditions, (5.48)

and

$$\langle X_i(t) \rangle^U \equiv \max_{\tilde{\mathbf{y}}(t)} \quad x_{0,i} + \sum_{r=1}^R s_{i,r} \tilde{y}_{\mathbf{e}_r}(t)$$

s.t. $\tilde{\mathbf{y}}(t)$ satisfies necessary (5.49)

moment conditions at time t,

respectively.
5.6.4 Necessary Conditions on the Moment Sequence

Third, we can derive necessary conditions that the moments \mathbf{y} must satisfy. Most of these conditions are directly analogous to the conditions appearing in Section (5.3.5). However, there are also a few subtle but important differences.

Dynamics-Based Necessary Conditions

The dynamics-based necessary conditions that we saw in Section 5.3.5 were derived from the CME (5.8) expressed in terms of the reduced states, $\hat{\mathbf{x}} \in \hat{\mathcal{X}}$. To derive analogous necessary conditions for the moments \mathbf{y} , we need an analogous CME expressed in terms of the states $\boldsymbol{\epsilon} \in \mathcal{E}$, which is

$$\frac{dP_{\mathcal{E}}}{dt}(\boldsymbol{\epsilon},t) = \sum_{r=1}^{R} \left[P_{\mathcal{E}}(\boldsymbol{\epsilon} - \mathbf{e}_r, t) \alpha_r(\boldsymbol{\epsilon} - \mathbf{e}_r) - P_{\mathcal{E}}(\boldsymbol{\epsilon}, t) \alpha_r(\boldsymbol{\epsilon}) \right], \quad \forall \boldsymbol{\epsilon} \in \mathcal{E},$$
(5.50)

where

$$\alpha_r(\boldsymbol{\epsilon}) \equiv a_r(\mathbf{x}_0 + \mathbf{S}\boldsymbol{\epsilon}). \tag{5.51}$$

From this CME, we can derive an analogue of Equation (5.12), describing the dynamics of the moments **y**:

$$\frac{d\mathbf{y}_L}{dt}(t) = \mathbf{A}_L^{\mathcal{E}} \mathbf{y}_L(t) + \mathbf{A}_H^{\mathcal{E}} \mathbf{y}_H(t).$$
(5.52)

The matrices $\mathbf{A}_{L}^{\mathcal{E}}$ and $\mathbf{A}_{H}^{\mathcal{E}}$ are clearly analogous to the matrices \mathbf{A}_{L} and \mathbf{A}_{H} appearing in Equation (5.12) and can be constructed in essentially the same way. Details can be found in Appendix D.

Now, in seeking the analogue of Necessary Condition (5.13), it is tempting to simply set the time derivatives on the left-hand side equal to zero, giving

$$\mathbf{0} = \mathbf{A}^{\mathcal{E}} \begin{bmatrix} \mathbf{y}_L(t) \\ \mathbf{y}_H(t) \end{bmatrix}, \qquad (5.53)$$

where $[\mathbf{A}_{L}^{\mathcal{E}} \ \mathbf{A}_{H}^{\mathcal{E}}] \equiv \mathbf{A}^{\mathcal{E}}$. This is, indeed, a necessary condition for a distribution

 $P_{\mathcal{E}}(\cdot, t)$ which is not changing with time. However, remember that our goal is to write a necessary condition for a time-invariant distribution $P(\cdot, t)$. It turns out that the time invariance of $P(\cdot, t)$ does not imply the time invariance of $P_{\mathcal{E}}(\cdot, t)$.

To see this, consider the simple reversible reaction system $A \Longrightarrow B$ with some nonzero initial count of A molecules and nonzero rate constants for both the forward and reverse reactions. In the steady-state distribution $P_{ss}(\cdot) = \lim_{t \to +\infty} P(\cdot, t)$, there will be some nonzero probability associated with every $\mathbf{x} \in \mathcal{X}$, and probability will be flowing between the reachable states in a way that the flow into each state balances the flow out. Probability can only flow from state to state if reactions are occurring. This implies that the state of the system, represented in terms of the reaction extents, $\boldsymbol{\epsilon}$ is constantly changing. In particular, both reaction extents are increasing with time at a constant rate, as the reactions continually fire in the dynamic equilibrium. In turn, this implies that the moments \mathbf{y} based on the reaction extents $\boldsymbol{\epsilon}$ are not constant with time. Thus, $\frac{d\mathbf{y}_L}{dt} = \mathbf{0}$ is not a necessary condition for a stochastic chemical kinetic system at steady state.

What do we use, then, for our necessary steady-state condition? While we cannot insist that $\frac{d\mathbf{y}_L}{dt} = \mathbf{0}$, it is still true that we must have $\frac{d\boldsymbol{\mu}_L}{dt} = \mathbf{0}$. Moreover, as we show in Appendix D, the time derivatives of the moments $\boldsymbol{\mu}_L$ can be related to those of \mathbf{y}_L through a linear equation:

$$\frac{d\boldsymbol{\mu}_L}{dt}(t) = \mathbf{C} \frac{d\mathbf{y}_L}{dt}(t).$$
(5.54)

Our steady-state necessary condition is then obtained by pre-multiplying both sides of Equation (5.53) by the matrix **C**:

$$\mathbf{0} = \mathbf{C}\mathbf{A}^{\mathcal{E}} \begin{bmatrix} \mathbf{y}_{L,\mathrm{ss}}(t) \\ \mathbf{y}_{H,\mathrm{ss}}(t) \end{bmatrix}.$$
 (5.55)

How about the necessary condition for the transient analysis? Fortunately, in this case, none of the above concerns apply, and the analogy is more straight-forward. We

simply set

$$\mathbf{y}_{L}(t) - e^{\rho t} \mathbf{y}_{L}(0) = (\mathbf{A}_{L}^{\mathcal{E}} - \rho \mathbf{I}) \mathbf{u}_{L}^{(\rho)} + \mathbf{A}_{H}^{\mathcal{E}} \mathbf{u}_{H}^{(\rho)}, \quad \forall \rho \in \mathcal{R},$$
(5.56)

where $\mathbf{u}^{(\rho)}$ is a vector of generalized moments defined by

$$u_{\mathbf{j}}^{(\rho)} \equiv \int_{0}^{t} e^{\rho(t-\tau)} y_{\mathbf{j}}(\tau) d\tau, \quad \forall \mathbf{j} \in \mathbb{N}^{R}, \forall \rho \in \mathcal{R}.$$
(5.57)

Probability-Based Necessary Conditions

The probability-based necessary conditions are directly analogous to those we had before. Since the total probability must always be one, we have

$$y_{\rm ss,0} = 1 \tag{5.58}$$

for the steady-state analysis and

$$y_{\mathbf{0}}(t) = 1 \tag{5.59}$$

and

$$u_{\mathbf{0}}^{(\rho)} = \begin{cases} T & \text{if } \rho = 0, \\ \frac{e^{\rho T} - 1}{\rho} & \text{otherwise,} \end{cases} \quad \forall \rho \in \mathcal{R},$$
(5.60)

for the dynamic analysis.

Furthermore, since

$$\sum_{\boldsymbol{\epsilon}\in\mathcal{E}} q^2(\boldsymbol{\epsilon}) \hat{P}_{\mathcal{E}}(\boldsymbol{\epsilon},t) \ge 0, \quad \forall t \ge 0,$$
(5.61)

for an arbitrary polynomial $q(\cdot)$, we have the LMIs

$$\mathbf{M}_{n}^{\mathbf{0}}(\mathbf{y}_{\mathrm{ss}}) \succeq \mathbf{0}, \tag{5.62}$$

for the steady-state analysis, and

$$\mathbf{M}_{n}^{\mathbf{0}}(\mathbf{y}(t)) \succeq \mathbf{0} \quad \text{and} \quad \mathbf{M}_{n}^{\mathbf{0}}(\mathbf{u}^{(\rho)}) \succeq \mathbf{0}, \ \forall \rho \in \mathcal{R}.$$
 (5.63)

for the transient analysis.

Support-Based Necessary Conditions

We've seen that the reachable set \mathcal{E} must be contained in the polyhedral superset $\overline{\mathcal{E}}$. This implies that $P_{\mathcal{E}}(\cdot, t)$ is supported on $\overline{\mathcal{E}}$. We then have a family of LMIs based on the inequality

$$\sum_{\epsilon \in \mathcal{E}} g(\epsilon) q^2(\epsilon) P_{\mathcal{E}}(\epsilon, t) \ge 0, \quad \forall t \ge 0,$$
(5.64)

where, again, $q(\cdot)$ is an arbitrary polynomial, and $g(\cdot)$ is one of the several polynomials defining $\overline{\mathcal{E}}$ through an inequality of the form $g(\epsilon) \ge 0$. In particular, polynomials $g(\cdot)$ of the form $g(\epsilon) = \epsilon_r$ for $r \in \{1, \ldots, R\}$ lead to the LMIs

$$\mathbf{M}_{n-1}^{\mathbf{e}_r}(\mathbf{y}_{\rm ss}) \succeq \mathbf{0}, \quad \forall r \in \{1, \dots, R\},$$
(5.65)

while polynomials $g(\cdot)$ of the form $g(\boldsymbol{\epsilon}) = x_{0,j} + \sum_{r=1}^{R} s_{j,r} \epsilon_r$ for $j \in \{1, \ldots, N\}$ lead to the LMIs

$$x_{0,j}\mathbf{M}_{n-1}^{\mathbf{0}}(\mathbf{y}_{ss}) + \sum_{r=1}^{R} s_{j,r}\mathbf{M}_{n-1}^{\mathbf{e}_{r}}(\mathbf{y}_{ss}) \succeq \mathbf{0}, \quad \forall j \in \{1, \dots, N\}.$$
 (5.66)

Of course, the same LMIs must be satisfied by the moment vectors $\mathbf{y}(t)$ and $\mathbf{u}^{(\rho)}$ for all $\rho \in \mathcal{R}$.

The LMIs described by Necessary Condition (5.66) reflect the fact that molecular counts must be nonnegative and are thus analogous to the LMIs described by Necessary Conditions (5.24) and (5.25) in Section 5.3.5. However, the LMIs described by Necessary Condition (5.65) have no analogue in Section 5.3.5

5.6.5 Semidefinite Programs

Using all of the necessary conditions appearing in Section 5.6.4 as constraints in the abstract optimization problems of Section 5.6.3, we obtain two SDPs: one for the steady-state problem,

$$\langle X_i \rangle_{\rm ss}^U \equiv \max_{\tilde{\mathbf{y}}} \quad x_{0,i} + \sum_{r=1}^R s_{i,r} \tilde{y}_{\mathbf{e}_r}$$
s.t. $\mathbf{C} \mathbf{A}^{\mathcal{E}} \begin{bmatrix} \tilde{\mathbf{y}}_L(t) \\ \tilde{\mathbf{y}}_H(t) \end{bmatrix} = \mathbf{0},$

$$\tilde{y}_{\mathbf{0}} = 1,$$

$$\mathbf{M}_n^{\mathbf{0}}(\tilde{\mathbf{y}}) \succeq \mathbf{0},$$

$$\mathbf{M}_{n-1}^{\mathbf{e}_r}(\tilde{\mathbf{y}}) \succeq \mathbf{0}, \quad \forall r \in \{1, \dots, R\},$$

$$x_{0,j} \mathbf{M}_{n-1}^{\mathbf{0}}(\tilde{\mathbf{y}}) + \sum_{r=1}^R s_{j,r} \mathbf{M}_{n-1}^{\mathbf{e}_j}(\tilde{\mathbf{y}}) \succeq \mathbf{0}, \quad \forall j \in \{1, \dots, N\},$$

and one for the transient problem,

$$\begin{split} \langle X_{i}(t) \rangle^{U} &\equiv \max_{\substack{\tilde{\mathbf{y}}(t)\\\tilde{\mathbf{u}}^{(\rho)},\forall \rho \in \mathcal{R}}} x_{0,i} \tilde{y}_{0}(t) + \sum_{r=1}^{R} s_{i,r} \tilde{y}_{e_{r}}(t) \\ \text{s.t.} \quad \tilde{\mathbf{y}}_{L}(t) - e^{\rho t} \tilde{\mathbf{y}}_{L}(0) = (\mathbf{A}_{L}^{\mathcal{E}} - \rho \mathbf{I}) \tilde{\mathbf{u}}_{L}^{(\rho)} + \mathbf{A}_{H}^{\mathcal{E}} \tilde{\mathbf{u}}_{H}^{(\rho)}, \quad \forall \rho \in \mathcal{R}, \\ \tilde{y}_{0}(t) = 1, \\ \tilde{u}_{0}^{(\rho)} &= \begin{cases} T & \text{if } \rho = 0, \\ \frac{e^{\rho T} - 1}{\rho} & \text{otherwise}, \end{cases} \quad \forall \rho \in \mathcal{R}, \\ \mathbf{M}_{n}^{\mathbf{0}}(\tilde{\mathbf{y}}(t)) \succeq \mathbf{0}, \\ \mathbf{M}_{n}^{\mathbf{0}}(\tilde{\mathbf{u}}^{(\rho)}) \succeq \mathbf{0}, \quad \forall r \in \{1, \dots, R\}, \\ \mathbf{M}_{n-1}^{\mathbf{e}_{r}}(\tilde{\mathbf{y}}(t)) \succeq \mathbf{0}, \quad \forall r \in \{1, \dots, R\}, \quad \forall \rho \in \mathcal{R}, \\ x_{0,j}\mathbf{M}_{n-1}^{\mathbf{0}}(\tilde{\mathbf{y}}(t)) + \sum_{r=1}^{R} s_{j,r}\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\tilde{\mathbf{y}}(t)) \succeq \mathbf{0}, \quad \forall j \in \{1, \dots, N\}, \\ x_{0,j}\mathbf{M}_{n-1}^{\mathbf{0}}(\tilde{\mathbf{u}}^{(\rho)}) + \sum_{r=1}^{R} s_{j,r}\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\tilde{\mathbf{u}}^{(\rho)}) \succeq \mathbf{0}, \quad \forall j \in \{1, \dots, N\}, \\ \forall \rho \in \mathcal{R}. \end{cases}$$

$$\end{split}$$

These SDPs have the same parameters as those described in Section 5.3.6.

Solving these particular SDPs gives an upper bound on the mean molecular count of species i at steady-state or at an arbitrary time $t \ge 0$, respectively. Of course, variations on this theme can be used to bound other statistical descriptions of the reaction system.

5.7 Revisiting the Pathological Example

In this section, we apply the alternative bounding method based on reaction extents (Section 5.6) to Reaction System (5.28), which we used in Section 5.4.2 to demonstrate a fundamental shortcoming of the bounding method based solely on independent species (Section 5.3). We will see that the alternative bounding method gives tighter

bounds for both the steady-state and dynamic analyses.

5.7.1 Steady-State Analysis

Solving SDP (5.67) with m = 2 to upper bound the steady-state mean molecular count of species B, we obtain $\langle B \rangle_{ss}^{U} = 3$. This upper bound is consistent with our observation in Section 5.4.2 that the stoichiometry and initial condition for Reaction System (5.28) imply that $x_B(t) \leq 3$ for all $t \geq 0$. Admittedly, it is, in some sense, still a trivial upper bound, as it is can also be calculated by a simple analysis of the stoichiometry. However, it is unmistakably a substantial improvement over the bound of $\langle B \rangle_{ss}^{U} = 6$ calculated in Section 5.4.2.

5.7.2 Transient Analysis

Repeating the transient analysis of Section (5.4.2), but using SDP (5.68) in place of SDP (5.27), we obtain the bounds shown in Figure 5-2. These bounds are tighter than those shown in Figure 5-1. This confirms our earlier suggestion that the over-approximation built into $\bar{\mathcal{X}}$ might be degrading the quality of the bounds produced by SDP (5.27).



Figure 5-2: Time-varying bounds on the mean molecular counts of species A, B, and D for Reaction System (5.28). These bounds were calculated by solving SDP (5.68) (and its minimization counterpart) over a range of t values, using m = 2 and $\mathcal{R} = \{0, -3, -8\}$. Compare with Figure 5-1.

5.8 Discussion

5.8.1 Weakness of the Steady-State Bounds

The fact that the bound $\langle B \rangle_{ss}^{U} = 3$ obtained in Section 5.7.1 is equal to the stoichiometric bound suggests that the steady-state SDP is not incorporating information about the relative time scales of the competing reactions. For the system to reach a steady-state in which $\langle B \rangle = 3$, all of the probability would have to flow from the initial state $\mathbf{x}(0) = (2, 3, 3, 0)$ via the second reaction, with no probability flowing via the first reaction. This could only occur if the second reaction was infinitely more likely than the first. But, whenever $x_B > 0$ (in particular, for the initial condition), this is not the case. Thus, the steady-state formulation that we have presented in this chapter seems to be missing an aspect of the dynamics inherent in the system. We will explore augmenting the steady-state formulation to add in this dynamic information in Chapter 6.

5.8.2 Strictly Better Bounds Are Not Guaranteed

While the alternative formulation of the bounding method described in Section 5.6 will always produce bounds that are at least as good as those produced by the original formulation described in Section (5.3), the bounds are not necessarily strictly better. In particular, if the set $\bar{\mathcal{X}}$ does not over-approximate the set $\hat{\mathcal{X}}$ in the sense described in Section 5.4.3, then the bounds produced by the two methods will be the same. An example where this occurs is given in Appendix D.

5.9 Conclusion

In this chapter, we have seen that the bounding method described in Chapters 3 and 4 fails to account for the irreversibility of reactions, and thus can perform poorly when irreversibility is important in defining the set of reachable states. In an effort to mitigate this problem, we described an alternative formulation of the bounding method which explicitly takes this irreversibility into account. We then demonstrated with an example that this alternative formulation can yield tighter bounds than the original formulation in both steady-state and transient analyses. Finally, we demonstrated that while the alternative formulation yields bounds that are at least as good as those given by the original formulation, the bounds are not necessarily strictly better.

Chapter 6

An Additional Steady-State Necessary Condition

The contents of this chapter have not yet been published.

6.1 Introduction

In this chapter, we derive an additional, relatively simple necessary condition on the low-order steady-state moments, $\mu_{L,ss}$. This necessary condition can be added as a constraint to any of the steady-state bounding SDPs that we have derived previously, potentially reducing the feasible set and tightening the bounds. We give a specific example of a system where this additional constraint leads to dramatically tighter bounds.

6.2 Deriving the Necessary Condition

By definition of the \mathbf{A}_L and \mathbf{A}_H matrices, we have

$$\frac{d\boldsymbol{\mu}_{L}}{dt}(t) = \mathbf{A}_{L}\boldsymbol{\mu}_{L}(t) + \mathbf{A}_{H}\boldsymbol{\mu}_{H}(t),$$

$$= \begin{bmatrix} \mathbf{A}_{L} & \mathbf{A}_{H} \end{bmatrix} \begin{bmatrix} \boldsymbol{\mu}_{L}(t) \\ \boldsymbol{\mu}_{H}(t) \end{bmatrix},$$

$$= \mathbf{A} \begin{bmatrix} \boldsymbol{\mu}_{L}(t) \\ \boldsymbol{\mu}_{H}(t) \end{bmatrix}.$$
(6.1)

Integrating both sides with respect to time, we obtain

$$\int_{0}^{T} \frac{d\boldsymbol{\mu}_{L}}{dt}(t)dt = \int_{0}^{T} \left(\mathbf{A} \begin{bmatrix} \boldsymbol{\mu}_{L}(t) \\ \boldsymbol{\mu}_{H}(t) \end{bmatrix} \right) dt,$$

$$\boldsymbol{\mu}_{L}(T) - \boldsymbol{\mu}_{L}(0) = \mathbf{A} \begin{bmatrix} \int_{0}^{T} \boldsymbol{\mu}_{L}(t)dt \\ \int_{0}^{T} \boldsymbol{\mu}_{H}(t)dt \end{bmatrix}.$$
(6.2)

It follows that

$$\boldsymbol{\mu}_L(T) - \boldsymbol{\mu}_L(0) \in \operatorname{rng}(\mathbf{A}), \quad \forall T \ge 0,$$
(6.3)

where $rng(\mathbf{A})$ denotes the range (i.e., column space) of the matrix \mathbf{A} . Assuming that the limit

$$\lim_{T \to +\infty} \boldsymbol{\mu}_L(T) \equiv \boldsymbol{\mu}_{L,\mathrm{ss}} \tag{6.4}$$

exists, it follows that

$$\boldsymbol{\mu}_{L,\mathrm{ss}} - \boldsymbol{\mu}_{L}(0) \in \mathrm{rng}(\mathbf{A}). \tag{6.5}$$

Proof. By definition of the limit, it follows that for every $\epsilon > 0$ there exists a $T_{\epsilon} > 0$ such that

$$||\boldsymbol{\mu}_L(T) - \boldsymbol{\mu}_{L,\mathrm{ss}}|| < \epsilon, \quad \forall T > T_\epsilon.$$
(6.6)

It follows that the sequence of points $\{\mu_L(n)\}_{n\in\mathbb{N}}$ converges to $\mu_{L,ss}$, because, for

every $\epsilon > 0$, there exists $N_{\epsilon} \equiv \lceil T_{\epsilon} \rceil \in \mathbb{N}$ such that

$$||\boldsymbol{\mu}_L(n) - \boldsymbol{\mu}_{L,\mathrm{ss}}|| < \epsilon, \quad \forall n \in \mathbb{N} \text{ s.t. } n > N_{\epsilon}.$$
(6.7)

From Statement (6.3), it follows that $\{\boldsymbol{\mu}_L(n)\}_{n\in\mathbb{N}}\subset \operatorname{rng}(\mathbf{A})$. Since $\operatorname{rng}(\mathbf{A})$ is a closed set, it follows that the every limit point of $\{\boldsymbol{\mu}_L(n)\}_{n\in\mathbb{N}}$ is an element of $\operatorname{rng}(\mathbf{A})$. In particular, $\boldsymbol{\mu}_{L,ss}$, as a limit point of $\{\boldsymbol{\mu}_L(n)\}_{n\in\mathbb{N}}$, is an element of $\operatorname{rng}(\mathbf{A})$. \Box

Statement (6.5) is equivalent to the statement that there exists some $\mathbf{w} \in \mathbb{R}^{\binom{M+\hat{N}}{M}}$ such that

$$\boldsymbol{\mu}_{L,\mathrm{ss}} - \boldsymbol{\mu}_{L}(0) = \mathbf{A}\mathbf{w}. \tag{6.8}$$

6.3 Using the Necessary Conditions as Constraints

We may now include this Necessary Condition (6.8) as a constraint in our steadystate SDPs, expressing it in terms of proxy decision variables. For example, adding this constraint to the SDP for calculating an upper bound on the steady-state mean molecular count of species i, we obtain

$$\langle X_i \rangle^U \equiv \max_{\tilde{\boldsymbol{\mu}}, \tilde{\mathbf{w}}} \quad v_{i,0} + \sum_{j=1}^{\hat{N}} v_{i,j} \tilde{\boldsymbol{\mu}}_{\mathbf{e}_j}$$
s.t. $\mathbf{A} \tilde{\boldsymbol{\mu}} = \mathbf{0},$

$$\tilde{\boldsymbol{\mu}}_L - \boldsymbol{\mu}_L(0) = \mathbf{A} \tilde{\mathbf{w}},$$

$$\tilde{\boldsymbol{\mu}}_0 = 1,$$

$$\mathbf{M}_n^0(\tilde{\boldsymbol{\mu}}) \succeq \mathbf{0},$$

$$v_{c,0} \mathbf{M}_{n-1}^0(\tilde{\boldsymbol{\mu}}) + \sum_{j=1}^{\hat{N}} v_{c,j} \mathbf{M}_{n-1}^{\mathbf{e}_j}(\tilde{\boldsymbol{\mu}}) \succeq \mathbf{0}, \quad \forall c \in \{1, \dots, N\}.$$

$$(6.9)$$

Interestingly, the additional constraint incorporates information about the initial moments (and thus the initial distirbution) which does not appear explicitly in our previous steady-state bounding SDPs. In these previous SDPs, the initial distribution was reflected in the invariant values. However, there are usually many different distributions (with different moments) corresponding to the same invariant values. The additional constraint effectively specifies which of these distributions is the true initial distribution.

6.3.1 Redundancy

Of course, the equation $\tilde{\boldsymbol{\mu}}_L - \boldsymbol{\mu}_L(0) = \mathbf{A}\tilde{\mathbf{w}}$ is really a system of equations, with one equation for each row of the **A** matrix. By it's construction, the first row of the **A** matrix is necessarily all zeros, implying the equation

$$\tilde{\mu}_0 = \mu_0(0). \tag{6.10}$$

Assuming that the initial distribution has been appropriately defined, we have $\mu_0(0) =$ 1. Thus, Equation (6.10) reduces to $\tilde{\mu}_0 = 1$, which we have already included as a constraint. It follows that the first row of $\tilde{\mu}_L - \mu_L(0) = \mathbf{A}\tilde{\mathbf{w}}$ is redundant and can be eliminated.

Let $\hat{\mathbf{A}}$ be equal to the \mathbf{A} matrix with the first row removed. Furthermore, let $\tilde{\boldsymbol{\mu}}_{1:m}$ denote a vector of the elements of $\tilde{\boldsymbol{\mu}}$ corresponding to multi-indices of order 1 through m, in graded descending order. Let $\boldsymbol{\mu}_{1:m}(0)$ be defined similarly. Then, we can express all of the equations implied by $\tilde{\boldsymbol{\mu}}_L - \boldsymbol{\mu}_L(0) = \mathbf{A}\tilde{\mathbf{w}}$ except the redundant equality by writing

$$\tilde{\boldsymbol{\mu}}_{1:m} - \boldsymbol{\mu}_{1:m}(0) = \hat{\mathbf{A}}\tilde{\mathbf{w}}$$
(6.11)

in its place. This gives us the equivalent SDP

$$\langle X_i \rangle^U \equiv \max_{\tilde{\boldsymbol{\mu}}, \tilde{\mathbf{w}}} \quad v_{i,0} + \sum_{j=1}^{\hat{N}} v_{i,j} \tilde{\mu}_{\mathbf{e}_j}$$
s.t. $\mathbf{A}\tilde{\boldsymbol{\mu}} = \mathbf{0},$

$$\tilde{\boldsymbol{\mu}}_{1:m} - \boldsymbol{\mu}_{1:m}(0) = \hat{\mathbf{A}}\tilde{\mathbf{w}},$$

$$\tilde{\boldsymbol{\mu}}_{\mathbf{0}} = 1,$$

$$\mathbf{M}_n^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}) \succeq \mathbf{0},$$

$$v_{c,0} \mathbf{M}_{n-1}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}) + \sum_{j=1}^{\hat{N}} v_{c,j} \mathbf{M}_{n-1}^{\mathbf{e}_j}(\tilde{\boldsymbol{\mu}}) \succeq \mathbf{0}, \quad \forall c \in \{1, \dots, N\}.$$

$$(6.12)$$

6.4 When Does it Help?

6.4.1 A Rank Condition

Now, the question is: under what circumstances will the additional constraint, $\tilde{\boldsymbol{\mu}}_{1:m} - \boldsymbol{\mu}_{1:m}(0) = \hat{\mathbf{A}}\tilde{\mathbf{w}}$, improve the quality of the bound, $\langle X_i \rangle^U$? We have added a constraint, so SDP (6.12) is potentially a restriction of the our typical mean-bounding SDP. However, if the columns of $\hat{\mathbf{A}}$ span the space $\mathbb{R}^{\binom{m+\hat{N}}{m}-1}$ in which $\tilde{\boldsymbol{\mu}}_{1:m} - \boldsymbol{\mu}_{1:m}(0)$ exists, then $\tilde{\boldsymbol{\mu}}_{1:m} - \boldsymbol{\mu}_{1:m}(0) = \hat{\mathbf{A}}\tilde{\mathbf{w}}$ isn't a restriction at all. It is only a restriction when the columns of $\hat{\mathbf{A}}$ do not span $\mathbb{R}^{\binom{m+\hat{N}}{m}-1}$. It other words, $\tilde{\boldsymbol{\mu}}_{1:m} - \boldsymbol{\mu}_{1:m}(0) = \hat{\mathbf{A}}\tilde{\mathbf{w}}$ is only a restriction when $\hat{\mathbf{A}}$ is not full row rank.

6.4.2 An Invariant Polynomial

What does it mean if \hat{A} is not full row rank? If \hat{A} is not full row rank, there exists some $\hat{q} \neq 0$ such that

$$\hat{\mathbf{q}}^{\mathrm{T}}\hat{\mathbf{A}} = \mathbf{0}^{\mathrm{T}}.\tag{6.13}$$

Letting $\mathbf{q} \equiv (0, \mathbf{q})$, it follows that $\mathbf{q} \neq \mathbf{0}$ and

$$\mathbf{q}^{\mathrm{T}}\mathbf{A} = \mathbf{0}^{\mathrm{T}}.$$
 (6.14)

Then, for moments $\boldsymbol{\mu}(t)$ of any distribution $\hat{P}(\cdot, t)$, we have

$$\mathbf{q}^{\mathrm{T}} \mathbf{A} \boldsymbol{\mu}(t) = 0, \quad \forall t \ge 0.$$
(6.15)

Using the definition of the A matrix, this implies

$$\mathbf{q}^{\mathrm{T}} \frac{d\boldsymbol{\mu}_L}{dt}(t) = \frac{d}{dt} \mathbf{q}^{\mathrm{T}} \boldsymbol{\mu}_L(t) = 0, \quad \forall t \ge 0.$$
(6.16)

In other words, a nontrivial linear combination of the low-order moments is constant with time.

This is equivalent to the statement that there exists some polynomial

$$q(\hat{\mathbf{x}}) = \sum_{|\mathbf{j}| \le m} q_{\mathbf{j}} \hat{\mathbf{x}}^{\mathbf{j}}$$
(6.17)

such that

$$\sum_{\hat{\mathbf{x}}\in\mathcal{X}} q(\hat{\mathbf{x}}) \frac{d\hat{P}}{dt}(\hat{\mathbf{x}}, t) = 0, \quad \forall t \ge 0.$$
(6.18)

6.5 Example

Consider the reaction system

$$A + B \xrightarrow{c_1} C$$

$$A + B \xrightarrow{c_2} D$$
(6.19)

with initial molecular counts of $A_0 = 4$, $B_0 = 3$, $C_0 = 0$, and $D_0 = 0$, and rate constants $c_1 = 1 \text{ s}^{-1}$ and $c_2 = 3 \text{ s}^{-1}$. Selecting species C and D as the independent species, and setting m = 2, we can solve our original SDP to calculate bounds on $\langle C \rangle_{ss}$. Doing so gives the stoichiometric bounds $\langle C \rangle_{ss}^L = 0$ and $\langle C \rangle_{ss}^U = 3$. While these bounds are valid, they are disappointingly loose.

On the other hand, if we solve SDP (6.12) for this system, we obtain $\langle C \rangle_{ss}^{L} = \langle C \rangle_{ss}^{U} = 0.75$, which implies that $\langle C \rangle_{ss} = 0.75$. Solving the CME for this system, we can confirm this value for $\langle C \rangle_{ss}$.

Why does the additional constraint make such a dramatic difference for this system? If we were to examine the propensity functions for the two reactions in the system, we would find that they have exactly the same functional form, and differ only by a scalar, $c_2/c_1 = 3$. This means that the average rate at which species D is formed is precisely 3 times the average rate at which species C is formed. Since we have zero molecules of both species initially, this implies that there are always, on average, 3 times as many D molecules as there are C molecules. This can be stated concisely as

$$3\langle C(t)\rangle = \langle D(t)\rangle, \quad \forall t \ge 0.$$
 (6.20)

Alternatively, it can be written in terms of moments:

$$3\mu_{(1,0)}(t) - 1\mu_{(0,1)}(t) = 0, \quad \forall t \ge 0,$$
(6.21)

which shows that we have a nontrivial linear combination of moments that is constant with time. This is consistent with the fact that $\hat{\mathbf{A}}$ matrix for this system,

$$\hat{\mathbf{A}} = \begin{bmatrix} 12 & -7 & -7 & 1 & 2 & 1 & 0 & 0 & 0 & 0 \\ 36 & -21 & -21 & 3 & 6 & 3 & 0 & 0 & 0 & 0 \\ 12 & 17 & -7 & -13 & -12 & 1 & 2 & 4 & 2 & 0 \\ 0 & 36 & 12 & -21 & -28 & -7 & 3 & 7 & 5 & 1 \\ 36 & -21 & 51 & 3 & -36 & -39 & 0 & 6 & 12 & 6 \end{bmatrix}$$
(6.22)

has a left null space which includes the vector $\hat{\mathbf{q}} = (3, -1, 0, 0, 0)$. Thus, this example is consistent with our general analysis in Section 6.4.

6.6 Extensions

The primary contribution of this chapter is the derivation of the necessary condition that there exists some $\mathbf{w} \in \mathbb{R}^{\binom{M+\hat{N}}{M}}$ such that

$$\boldsymbol{\mu}_{L,\mathrm{ss}} - \boldsymbol{\mu}_L(0) = \mathbf{A}\mathbf{w}. \tag{6.23}$$

However, this idea can be taken further.

For example, since each $\mu_{\mathbf{j}}(t) \ge 0$ for all t, it is relatively easy to show that the \mathbf{w} must be nonnegative. The argument is essentially the same as the one given above. The only real difference is that we use the fact that

$$\{\mathbf{A}\mathbf{w}:\mathbf{w}\ge0\}\tag{6.24}$$

is a closed set. This implies that we can also include the constraint $\tilde{\mathbf{w}} \geq \mathbf{0}$ in the bounding SDP, with potentially dramatic consequences.

In fact, for any polyhedral set \mathcal{P} , the set

$$\{\mathbf{A}\mathbf{w}:\mathbf{w}\in\mathcal{P}\}\tag{6.25}$$

is closed. We may be able to use this fact to obtain even tighter bounds, if we can show that

$$\int_0^T \boldsymbol{\mu}(t) dt \in \mathcal{P}, \quad \forall T \ge 0.$$
(6.26)

Taking the idea on step further still, since the vector $\int_0^T \boldsymbol{\mu}(t) dt$ satisfies the LMIs appearing in (6.9) for all times $T \geq 0$, we may be able to show that \mathbf{w} must satisfy these LMIs as well. However, the proof of this claim (if it exists) is less-straightforward, because the image of a set defined by LMIs under a linear mapping is not necessarily closed.

These ideas are fertile ground for future research. They have not yet been adequately explored.

6.7 Conclusion

We've seen that inclusion of a constraint based on Equation (6.8) in the steady-state analysis can yield significantly improved bounds. However, this constraint will only benefit the bounds in the situation where there is some nontrivial linear combination of the moments which is constant with time. We've seen one example where this is the case: when the propensity functions of two different reactions in the system differ by a constant factor. There may be other structural features of a reaction system for which the additional constraint benefits the bounds. Exploring these additional features is left for future work. Moreover, there are clear extensions of the main idea of this chapter, which are also left for future work.

Chapter 7

Improved Numerical Performance Through Linear Programming Approximations

The contents of this chapter have not yet been published. The ideas contained herein were developed in collaboration with Tim Varelmann, as part of his Master's Thesis.

7.1 Introduction

In the previous chapters, we noted that the SDPs we derived often cause numerical difficulties for off-the-shelf solvers like SeDuMi and Mosek. This limits the practical value of the SDPs we have derived. So, the aim of this chapter is to devise a strategy for solving these SDPs more reliably, or at least reliably outer-bounding their solutions.

7.2 Notation

In this section, we define some notation which is necessary to concisely express the ideas that will developed later in this chapter.

7.2.1 Invariants and Reduced Reachable States

In the previous chapters, we made use of the fact that reaction invariants can lead to a lower-dimensional representation of the state. In particular, the stoichiometry matrix **S** often has a nontrivial left null space, spanned by a set of basis vectors $\{\mathbf{b}_1, \ldots, \mathbf{b}_L\} \subset \mathbb{R}^N$. Setting

$$\mathbf{B} = \begin{bmatrix} \mathbf{b}_1^{\mathrm{T}} \\ \vdots \\ \mathbf{b}_L^{\mathrm{T}} \end{bmatrix}, \qquad (7.1)$$

we noted that the the state $\mathbf{X}(t)$ must satisfy the affine equation

$$\mathbf{BX}(t) = \mathbf{Bx}_0 \equiv \mathbf{f} \tag{7.2}$$

for all times $t \ge 0$. We referred to the vector **f** as the *invariant values*.

Since the *L* rows of **B** are linearly independent, it follows that **B** has *L* linearly independent columns, corresponding to *L* chemical species. Let the selection matrix $\tilde{\mathbf{Q}} \in \{0,1\}^{N \times L}$ be defined such that the product $\mathbf{B}\tilde{\mathbf{Q}} \equiv \tilde{\mathbf{B}} \in \mathbb{R}^{L \times L}$ is a nonsingular matrix consisting of just these *L* columns. The molecular counts of the corresponding *L* species are given by $\tilde{\mathbf{Q}}^{\mathrm{T}}\mathbf{X}(t) \equiv \tilde{\mathbf{X}}(t)$. Let $\hat{\mathbf{Q}} \in \{0,1\}^{N \times (N-L)}$ be a complementary selection matrix such that $\hat{\mathbf{B}} \equiv \mathbf{B}\hat{\mathbf{Q}} \in \mathbb{R}^{L \times (N-L)}$ contains the $N - L \equiv \hat{N}$ columns of **B** not found in $\tilde{\mathbf{B}}$. It follows that $\hat{\mathbf{Q}}^{\mathrm{T}}\mathbf{X}(t) \equiv \hat{\mathbf{X}}(t)$ contains the molecular counts of the chemical species not represented in $\tilde{\mathbf{X}}(t)$, and that

$$\mathbf{Q} \equiv \begin{bmatrix} \hat{\mathbf{Q}} & \tilde{\mathbf{Q}} \end{bmatrix} \in \{0, 1\}^{N \times N}$$
(7.3)

is a permutation matrix, satisfying $\mathbf{Q}\mathbf{Q}^{\mathrm{T}} = \mathbf{I} \in \mathbb{R}^{N \times N}$. From this fact, it follows that

$$\mathbf{X}(t) = \mathbf{Q}\mathbf{Q}^{\mathrm{T}}\mathbf{X}(t) = \mathbf{Q}\begin{bmatrix} \hat{\mathbf{Q}}^{\mathrm{T}} \\ \tilde{\mathbf{Q}}^{\mathrm{T}} \end{bmatrix} \mathbf{X}(t) = \mathbf{Q}\begin{bmatrix} \hat{\mathbf{Q}}^{\mathrm{T}}\mathbf{X}(t) \\ \tilde{\mathbf{Q}}^{\mathrm{T}}\mathbf{X}(t) \end{bmatrix} = \mathbf{Q}\begin{bmatrix} \hat{\mathbf{X}}(t) \\ \tilde{\mathbf{X}}(t) \end{bmatrix}.$$
 (7.4)

Furthermore, we can rewrite Equation (7.2) as

$$\mathbf{f} = \mathbf{B}\mathbf{X}(t),$$

$$= \mathbf{B}\mathbf{Q}\mathbf{Q}^{\mathrm{T}}\mathbf{X}(t),$$

$$= \mathbf{B}\begin{bmatrix} \hat{\mathbf{Q}} & \tilde{\mathbf{Q}} \end{bmatrix} \begin{bmatrix} \hat{\mathbf{Q}}^{\mathrm{T}} \\ \tilde{\mathbf{Q}}^{\mathrm{T}} \end{bmatrix} \mathbf{X}(t),$$

$$= \begin{bmatrix} \hat{\mathbf{B}} & \tilde{\mathbf{B}} \end{bmatrix} \begin{bmatrix} \hat{\mathbf{X}}(t) \\ \tilde{\mathbf{X}}(t) \end{bmatrix},$$

$$= \hat{\mathbf{B}}\hat{\mathbf{X}}(t) + \tilde{\mathbf{B}}\tilde{\mathbf{X}}(t).$$
(7.5)

Rearranging (7.5), we obtain

$$\tilde{\mathbf{X}}(t) = \tilde{\mathbf{B}}^{-1}\mathbf{f} - \tilde{\mathbf{B}}^{-1}\hat{\mathbf{B}}\hat{\mathbf{X}}(t), \qquad (7.6)$$

which shows that specifying $\hat{\mathbf{X}}(t)$ determines the values of $\tilde{\mathbf{X}}(t)$. Accordingly, we labeled the species represented by $\hat{\mathbf{X}}(t)$ as the *independent species*, and those represented by $\tilde{\mathbf{X}}(t)$ as the *dependent species*. Let $I \subset \{1, \ldots, N\}$ denote the set of indices corresponding to the independent species. Similarly, let its complement $D \equiv I^C$ denote the set of indices corresponding to the dependent species. Since specifying $\hat{\mathbf{X}}(t)$ determines $\tilde{\mathbf{X}}(t)$, we saw that the full state of the system $\mathbf{X}(t)$ can be expressed entirely in terms of the lower-dimensional $\hat{\mathbf{X}}(t)$. In particular, from Equations (7.4) and (7.6), we have

$$\mathbf{X}(t) = \mathbf{Q} \begin{bmatrix} \hat{\mathbf{X}}(t) \\ \tilde{\mathbf{B}}^{-1}\mathbf{f} - \tilde{\mathbf{B}}^{-1}\hat{\mathbf{B}}\hat{\mathbf{X}}(t) \end{bmatrix} = \mathbf{Q} \left(\begin{bmatrix} \mathbf{0} \\ \tilde{\mathbf{B}}^{-1}\mathbf{f} \end{bmatrix} + \begin{bmatrix} \mathbf{I} \\ -\tilde{\mathbf{B}}^{-1}\hat{\mathbf{B}} \end{bmatrix} \hat{\mathbf{X}}(t) \right). \quad (7.7)$$

It follows that every full-dimensional reachable state $\mathbf{x} \in \mathcal{X} \subset \mathbb{N}^N$ can be expressed equivalently in terms of a *reduced reachable state* $\hat{\mathbf{Q}}^{\mathrm{T}}\mathbf{x} \equiv \hat{\mathbf{x}} \in \mathbb{N}^{\hat{N}}$. We denoted the set of reduced reachable states using the symbol $\hat{\mathcal{X}} \subset \mathbb{N}^{\hat{N}}$. Setting

$$\mathbf{v}_0 \equiv \mathbf{Q} \begin{bmatrix} \mathbf{0} \\ \tilde{\mathbf{B}}^{-1} \mathbf{f} \end{bmatrix}$$
(7.8)

and

$$\mathbf{V} \equiv \mathbf{Q} \begin{bmatrix} \mathbf{I} \\ -\tilde{\mathbf{B}}^{-1}\hat{\mathbf{B}} \end{bmatrix},\tag{7.9}$$

Equation (7.7) can be written concisely as

$$\mathbf{X}(t) = \mathbf{v}_0 + \mathbf{V}\hat{\mathbf{X}}(t). \tag{7.10}$$

Since molecular counts must always be nonnegative, it follows that $\hat{\mathbf{X}}(t)$ must always be contained in the set

$$\bar{\mathcal{X}} = \left\{ \hat{\mathbf{x}} \in \mathbb{R}^{\hat{N}} : \mathbf{v}_0 + \mathbf{V}\hat{\mathbf{x}} \ge \mathbf{0} \right\}.$$
(7.11)

In what follows, we will let $v_{c,0}$ represent the *c*th component of the vector \mathbf{v}_0 . Similarly, we let $v_{c,j}$ represent the element in the *c*th row and *j*th column of the matrix \mathbf{V} .

7.3 Numerical Stability through an LP Relaxation

In this section, we describe a linear programming (LP) relaxation of a representative moment-based SDP. This relaxation, and the refinement described in Section 7.4, are based heavily on the work of Ahmadi et al. [1].

7.3.1 Relaxation Yields Valid Bounds

In all of our work thus far, we have been solving SDPs to calculate a bound on some actual quantity. For example, to calculate an upper bound on the steady-state mean molecular count of species i, we proposed solving the an SDP which can be expressed

concisely as:

$$p^* \equiv \max_{\tilde{\boldsymbol{\mu}}} \quad v_{i,0} + \sum_{j=1}^{\hat{N}} v_{i,j} \tilde{\boldsymbol{\mu}}_{\mathbf{e}_j}$$

s.t. $\mathbf{A}\tilde{\boldsymbol{\mu}} = \mathbf{0},$
 $\tilde{\boldsymbol{\mu}}_{\mathbf{0}} = 1,$ (7.12)
 $\mathbf{M}_n^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}) \succeq \mathbf{0},$
 $v_{c,0} \mathbf{M}_{n-1}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}) + \sum_{j=1}^{\hat{N}} v_{c,j} \mathbf{M}_{n-1}^{\mathbf{e}_j}(\tilde{\boldsymbol{\mu}}) \succeq \mathbf{0}, \quad \forall c \in \{1, \dots, N\},$

where the constraints are necessary conditions on the (unknown) true steady-state moment sequence μ_{ss} . We showed that p^* provides an upper bound on $\langle X_i \rangle_{ss}$, the essential points of the argument being:

- 1. Since the constraints of SDP (7.12) are necessary conditions for μ_{ss} , the point $\tilde{\mu} = \mu_{ss}$ is feasible.
- 2. The objective value corresponding to $\tilde{\mu} = \mu_{\rm ss}$ is

$$v_{i,0} + \sum_{j=1}^{\hat{N}} v_{i,j} \mu_{\mathbf{e}_j} = \langle X_i \rangle_{\mathrm{ss}}.$$
 (7.13)

Notice that this argument still holds if we relax any of the constraints in SDP (7.12). Relaxing any of the constraints will give a different problem, with an optimal value \tilde{p} which may be strictly greater than p^* . However, we are still guaranteed that $\tilde{p} \geq \langle X_i \rangle_{ss}$. The question then is this: can we relax SDP (7.12) in such a way that the problem becomes easier to solve, without introducing too much conservatism into the bound on $\langle X_i \rangle_{ss}$?

7.3.2 Relaxing the LMIs

Recall that the LMI $\mathbf{M}_n^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}) \succeq \mathbf{0}$ is equivalent to the condition

$$\mathbf{p}^{\mathrm{T}}\mathbf{M}_{n}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}})\mathbf{p} \geq 0, \quad \forall \mathbf{p} \in \mathbb{R}^{\binom{n+\tilde{N}}{n}}.$$
 (7.14)

Importantly, the inequality associated with each $\mathbf{p} \in \mathbb{R}^{\binom{n+\hat{N}}{n}}$ is linear with respect to the vector $\tilde{\boldsymbol{\mu}}$. If we modify SDP (7.12) by replacing the LMI $\mathbf{M}_n^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}) \succeq \mathbf{0}$ with Condition (7.14), and we replace the other N LMIs by analogous conditions, we obtain an equivalent optimization problem:

$$p^{*} = \max_{\tilde{\boldsymbol{\mu}}} \quad v_{i,0} + \sum_{j=1}^{\hat{N}} v_{i,j} \tilde{\boldsymbol{\mu}}_{\mathbf{e}_{j}}$$

s.t. $\mathbf{A}\tilde{\boldsymbol{\mu}} = \mathbf{0},$
 $\tilde{\boldsymbol{\mu}}_{\mathbf{0}} = 1,$
 $\mathbf{p}^{\mathrm{T}} \mathbf{M}_{n}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}) \mathbf{p} \ge 0, \quad \forall \mathbf{p} \in \mathbb{R}^{\binom{n+\hat{N}}{n}},$
 $\mathbf{n}^{\mathrm{T}} \left(v_{c,0} \mathbf{M}_{n-1}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}) + \sum_{j=1}^{\hat{N}} v_{c,j} \mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\tilde{\boldsymbol{\mu}}) \right) \mathbf{n} \ge 0, \quad \forall \mathbf{n} \in \mathbb{R}^{\binom{n-1+\hat{N}}{n-1}},$
 $\forall c \in \{1, \dots, N\}.$ (7.15)

This is a *semi-infinite linear program* – a linear program with infinitely many inequality constraints.

Suppose that we pick some finite set of vectors $\mathcal{P} \equiv \{\mathbf{p}_1, \dots, \mathbf{p}_{|\mathcal{P}|}\} \subset \mathbb{R}^{\binom{n+\hat{N}}{n}}$ and replace the semi-infinite constraint

$$\mathbf{p}^{\mathrm{T}}\mathbf{M}_{n}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}})\mathbf{p} \geq 0, \quad \forall \mathbf{p} \in \mathbb{R}^{\binom{n+\hat{N}}{n}}$$
(7.16)

with finitely many linear constraints

$$\mathbf{p}^{\mathrm{T}}\mathbf{M}_{n}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}})\mathbf{p} \geq 0, \quad \forall \mathbf{p} \in \mathcal{P}.$$
(7.17)

Suppose we similarly pick finite sets of vectors $\mathcal{N}_c \subset \mathbb{R}^{\binom{n-1+\hat{N}}{n-1}}$ for $c \in \{1, \ldots, N\}$, and replace the last N semi-infinite constraints with finitely many constraints. Doing so

results in the linear program:

$$\tilde{p} \equiv \max_{\tilde{\boldsymbol{\mu}}} \quad v_{i,0} + \sum_{j=1}^{\hat{N}} v_{i,j} \tilde{\boldsymbol{\mu}}_{\mathbf{e}_j}$$
s.t. $\mathbf{A}\tilde{\boldsymbol{\mu}} = \mathbf{0},$
 $\tilde{\mu}_{\mathbf{0}} = 1,$
 $\mathbf{p}^{\mathrm{T}} \mathbf{M}_n^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}) \mathbf{p} \ge 0, \quad \forall \mathbf{p} \in \mathcal{P},$
 $\mathbf{n}^{\mathrm{T}} \left(v_{c,0} \mathbf{M}_{n-1}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}) + \sum_{j=1}^{\hat{N}} v_{c,j} \mathbf{M}_{n-1}^{\mathbf{e}_j}(\tilde{\boldsymbol{\mu}}) \right) \mathbf{n} \ge 0, \quad \forall \mathbf{n} \in \mathcal{N}_c, \quad \forall c \in \{1, \dots, N\}.$
(7.18)

Clearly, this linear program is a relaxation of the semi-infinite program (7.15). It follows that $\tilde{p} \ge p^*$. It may be the case that $\tilde{p} - p^* > 0$, in which case solving LP (7.18) provides a strictly weaker upper bound on $\langle X_i \rangle_{ss}$ than that provided by solving SDP (7.12). However, if we cannot solve SDP (7.12) due to numerical difficulties, this gap in the bounds is somewhat irrelevant.

Given that we are focused on overcoming numerical limitations, solving an LP to obtain the bound is appealing. Linear programming is a more mature technology than semidefinite programming. In particular, linear programming solvers such as Gurobi and CPLEX are more numerically robust than their semidefinite programming counterparts. This gives us hope that in situations where we cannot solve SDP (7.12) due to numerical difficulties, we may still be able to solve LP (7.18) and obtain nontrivial bounds on $\langle X_i \rangle_{ss}$.

7.3.3 Choosing \mathcal{P} and \mathcal{N}_c

But the obvious question is: how do we choose the sets \mathcal{P} and \mathcal{N}_c for $c \in \{1, \ldots, N\}$? Generally speaking, the fewer vectors that we include in each set, the greater the relaxation and the greater the gap between \tilde{p} and p^* . This suggests that we should include a large number of vectors in each set, perhaps sampling uniformly from the unit sphere. However, there is a balance to be struck, because more parameter vectors imply more linear inequality constraints in the LP.

Ahmadi Vectors

In an effort to achieve this balance, Ahmadi et al. [1] suggest that an LMI featuring an $n \times n$ dimensional matrix **Y** should be replaced with n^2 linear inequalities, where the vectors defining these inequalities have at most two nonzero components, each equal to ± 1 . We will denote this set of vectors as \mathcal{A}_n . For example, in the case of n = 2, the set of vectors is,

$$\mathcal{A}_2 = \left\{ \begin{bmatrix} 1\\0 \end{bmatrix}, \begin{bmatrix} 0\\1 \end{bmatrix}, \begin{bmatrix} 1\\1 \end{bmatrix}, \begin{bmatrix} 1\\1 \end{bmatrix}, \begin{bmatrix} 1\\-1 \end{bmatrix} \right\}.$$
(7.19)

Similarly, in the case of n = 3, the set of vectors is,

$$\mathcal{A}_{3} = \left\{ \begin{bmatrix} 1\\0\\0 \end{bmatrix}, \begin{bmatrix} 0\\1\\0 \end{bmatrix}, \begin{bmatrix} 0\\0\\1 \end{bmatrix}, \begin{bmatrix} 1\\1\\0 \end{bmatrix}, \begin{bmatrix} 1\\1\\0 \end{bmatrix}, \begin{bmatrix} 1\\-1\\0 \end{bmatrix}, \begin{bmatrix} 1\\0\\1 \end{bmatrix}, \begin{bmatrix} 1\\0\\-1 \end{bmatrix}, \begin{bmatrix} 0\\1\\1 \end{bmatrix}, \begin{bmatrix} 0\\1\\1 \end{bmatrix}, \begin{bmatrix} 0\\1\\-1 \end{bmatrix} \right\}.$$
(7.20)

Notice that for each vector $\mathbf{a} \in \mathcal{A}_n$, there is no need for the vector $-\mathbf{a}$ to also be in \mathcal{A}_n , because the linear inequalities

$$\mathbf{a}^{\mathrm{T}}\mathbf{Y}\mathbf{a} \ge 0 \tag{7.21}$$

and

$$(-\mathbf{a})^{\mathrm{T}}\mathbf{Y}(-\mathbf{a}) \ge 0 \tag{7.22}$$

are equivalent.

If we set

$$\mathcal{P} \equiv \mathcal{A}_{\binom{n+\hat{N}}{n}} \tag{7.23}$$

and

$$\mathcal{N}_c \equiv \mathcal{A}_{\binom{n-1+\hat{N}}{n-1}}, \quad \forall c \in \{1, \dots, N\},$$
(7.24)

LP (7.18) is well-defined, and we can solve solve it to obtain an upper bound on $\langle X_i \rangle_{ss}$. Of course, solving the corresponding minimization problem, we can also calculate a lower bound.

A Bounded LP

One of the advantages of using the Ahmadi vectors is that they ensure the LP (7.18) is bounded, provided the set $\bar{\mathcal{X}}$ defined in Equation (7.11) is bounded. The reason for this is that $\mathcal{A}_{\binom{n-1+\hat{N}}{n-1}}$ always includes the vector $\mathbf{e}_1 \in \mathbb{R}^{\binom{n-1+\hat{N}}{n-1}}$. When this vector is included in the set \mathcal{N}_c , LP (7.18) will be constrained by the corresponding linear inequalities

$$\mathbf{e}_{1}^{\mathrm{T}}\left(v_{c,0}\mathbf{M}_{n-1}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}) + \sum_{j=1}^{\hat{N}} v_{c,j}\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\tilde{\boldsymbol{\mu}})\right)\mathbf{e}_{1} \ge 0, \quad \forall c \in \{1,\ldots,N\}.$$
(7.25)

Using the definition of the $\mathbf{M}_{n-1}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}})$ and $\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\tilde{\boldsymbol{\mu}})$ matrices given in Appendix B, we can show that these are equivalent to the linear inequalities

$$v_{c,0} + \sum_{j=1}^{\hat{N}} v_{c,j} \tilde{\mu}_{\mathbf{e}_j} \ge 0, \quad \forall c \in \{1, \dots, N\}.$$
 (7.26)

Consider an LP of the form

$$\max_{\hat{\mathbf{x}} \in \mathbb{R}^{\hat{N}}} \quad v_{i,0} + \sum_{j=1}^{\hat{N}} v_{i,j} \hat{x}_j$$

s.t. $v_{c,0} + \sum_{j=1}^{\hat{N}} v_{c,j} \hat{x}_j \ge 0, \quad \forall c \in \{1, \dots, N\}.$ (7.27)

We see that this LP maximizes the molecular count of species i over the set $\bar{\mathcal{X}}$. Accordingly, if the set $\bar{\mathcal{X}}$ is bounded, then the optimal value of LP (7.27) is bounded.

Now, while the decision variables differ between LP (7.18) and LP (7.27), the objective functions have the same form. Furthermore, the linear inequalities described by (7.26) are the same form as the constraints in (7.27). Accordingly, LP (7.18), with the linear inequalities described by (7.26) and other inequality and equality constraints, can be viewed as a *restriction* of LP (7.27). Since the objective value of LP (7.27) is bounded, it follows that the objective value of LP (7.18) is bounded.

7.3.4 Small Example

To demonstrate the use of LP (7.18), we now apply it in a numerical example.

Consider a variant of the Michaelis-Menten reaction system, described by Smadbeck and Kaznessis [64]:

$$S + E \xrightarrow{c_1} S:E$$

$$S:E \xrightarrow{c_3} P + E$$

$$P \xrightarrow{c_4} S$$

$$(7.28)$$

with initial molecular counts S = 10, E = 10, S:E = 0, and P = 0. In Chapter 3, we showed that we could obtain very tight bounds on $\langle S \rangle_{ss}$ and $\langle E \rangle_{ss}$ for a range of rate constant values by repeatedly solving SDP (7.12). These bounds are reproduced in the top panel of Figure 7-1.

The bottom panel of the figure shows analogous bounds produced by repeatedly solving LP (7.18), with the sets \mathcal{P} and \mathcal{N}_c for $c \in \{1, \ldots, N\}$ defined by Equations (7.23) and (7.24), respectively.

Clearly, for this example, the SDP-based bounds are superior. However, the LPbased bounds aren't terrible. In fact, for some values of the rate constant c_1 , the LP bounds are just as tight as the SDP bounds.



Figure 7-1: Bounds on the Michaelis-Menten system at steady state. Upper bounds are shown with open circles, while lower bounds are shown with stars. The top panel shows bounds produced by repeatedly solving SDP (7.12) for a range of values of the rate constant c_1 (the other rate constants being fixed at $c_2 = c_3 = c_4 = 1 \text{ s}^{-1}$). Similarly, the bottom panel shows bounds produced by repeatedly solving LP (7.18), with the sets \mathcal{P} and \mathcal{N}_c for $c \in \{1, \ldots, N\}$ defined by Equations (7.23) and (7.24), respectively. To provide a reference, both panels show the true means calculated by directly solving the CME. These CME solutions are shown as dashed lines. The solid lines connecting the circles and stars are not theoretically guaranteed bounds; they are included just to lead the eye.

7.3.5 Large Example

Next we apply LP (7.18) to a larger reaction system from Chapters 3 and 4, reproduced in Figure 7-2.



Figure 7-2: A larger reaction system

The rate constants are given in Appendix B. With an initial state consisting of 100 molecules of species A and F and 0 molecules of all other species, this system has over 74 billion reachable states. Thus, solving the CME for this system is impractical. In Chapter ??, we calculated bounds on $\langle A \rangle_{ss}$ for a range of rate constants by repeatedly solving SDP (7.12). These bounds are reproduced in the top panel of Figure 7-3.



Figure 7-3: A larger reaction system

The bottom panel of the figure shows analogous bounds produced by repeatedly solving LP (7.18), with the sets \mathcal{P} and \mathcal{N}_c for $c \in \{1, \ldots, N\}$ defined by Equations (7.23) and (7.24), respectively.

As in the Michaelis-Menten example, the SDP-based bounds are clearly superior. The LP-based bounds, while nontrivial, are much looser, with very wide gaps in the bounds for some rate constant values.

7.3.6 Summary

The LP-based bounds that we have described in Section 7.3 are robust in the sense that they can be obtained without excessive numerical issues. However, as demonstrated by the example in Section 7.3.5, these LP-based bounds are sometimes unacceptably loose.

7.4 Iterative Refinement of LP Bounds

In this section, we describe an algorithm for iteratively refining the bounds produced by solving LP (7.18). This is essentially a cutting plane algorithm, inspired by the work of Krishnan and Mitchell [38] and Ahmadi et al. [1].

7.4.1 Calculation of Cutting Planes

Suppose that we have solved LP (7.18), with the sets \mathcal{P} and \mathcal{N}_c for $c \in \{1, \ldots, N\}$ defined by Equations (7.23) and (7.24), respectively, and obtained an optimal solution $\tilde{\mu}^*$. Suppose we then calculate the eigenvalues of the matrix $\mathbf{M}_n^0(\tilde{\mu}^*)$ and find that one of these eigenvalues, λ , is strictly negative. This implies that $\tilde{\mu}^*$ is infeasible for SDP (7.12). Accordingly, we'd like to exclude it from the feasible set of LP (7.18). Let $\mathbf{p}_{\lambda} \in \mathbb{R}^{\binom{n+\hat{N}}{n}}$ be the normalized eigenvector corresponding to the eigenvalue λ . If we add the vector \mathbf{p}_{λ} to the set \mathcal{P} , we add the constraint

$$\mathbf{p}_{\lambda}^{\mathrm{T}} \mathbf{M}_{n}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}) \mathbf{p}_{\lambda} \ge 0 \tag{7.29}$$

to LP (7.18). This constraint renders the vector $\tilde{\mu}^*$ infeasible, because

$$\mathbf{p}_{\lambda}^{\mathrm{T}} \mathbf{M}_{n}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}^{*}) \mathbf{p}_{\lambda} = \mathbf{p}_{\lambda}^{\mathrm{T}} \lambda \mathbf{p}_{\lambda} = \lambda \mathbf{p}_{\lambda}^{\mathrm{T}} \mathbf{p}_{\lambda} = \lambda < 0.$$
(7.30)

The constraint $\mathbf{p}_{\lambda}^{\mathrm{T}} \mathbf{M}_{n}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}) \mathbf{p}_{\lambda} \geq 0$ is a cutting plane in the sense that it is guaranteed to exclude $\tilde{\boldsymbol{\mu}}^{*}$ from the feasible set of LP (7.18). At the same time, by the inequality's structure, we are guaranteed that it does not exclude any point in the feasible set of SDP (7.12).

Of course, if any of the matrices $v_{c,0}\mathbf{M}_{n-1}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}) + \sum_{j=1}^{\hat{N}} v_{c,j}\mathbf{M}_{n-1}^{\mathbf{e}_j}(\tilde{\boldsymbol{\mu}})$ for $c \in \{1, \ldots, N\}$ has a strictly negative eigenvalue, we could generate another cutting plane in much

the same way.

Having added the cutting planes to LP (7.18), we could again solve the problem, obtain a new optimal solution, add new cutting planes – and repeat the cycle indefinitely. As we add cutting planes, we systematically refine the feasible set of LP (7.18)so that it better and better approximates the feasible set of SDP (7.12). If we could refine the feasible set of LP (7.18) to the point where it was the same as the feasible set of SDP (7.12), the optimal values would also be the same.

7.4.2 A Cutting Plane Algorithm

The observations of the foregoing section suggest a procedure for iteratively refining the bounds produced by solving LP (7.18), which we formalize in Algorithm 4.

Algorithm 4 Iterative Refinement of LP Bounds through Cutting Planes Input:

- 1. The index $i \in \{1, ..., N\}$ corresponding to the species whose mean you'd like to bound.
- 2. The matrix \mathbf{A} .
- 3. The vectors $\{\mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_{\hat{N}}\}$
- 4. The maximum number of cutting plane iterations, k_{max} .

Output: An upper bound $\langle X_i \rangle_{ss}^U$ on the steady-state mean molecular count of species *i*.

Algorithm:

Set $\mathcal{P} := \mathcal{A}_{\binom{n+\hat{N}}{n}}$ and $\mathcal{N}_c := \mathcal{A}_{\binom{n-1+\hat{N}}{n-1}}$ for all $c \in \{1, \ldots, N\}$. Solve LP (7.18) and extract the optimal solution $\tilde{\mu}_0^*$ and the optimal value \tilde{p}_0 .

for $k = 1, \ldots, k_{\max} \operatorname{do}$

Calculate $\lambda := \lambda_{\min}(\mathbf{M}_n^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}_{k-1}^*))$ and the corresponding normalized eigenvector \mathbf{p}_{λ} .

if $\lambda < 0$ then

Set $\mathcal{P} := \mathcal{P} \cup \{\mathbf{p}_{\lambda}\}.$

end if

for $c = 1, \ldots, N$ do

Calculate $\lambda := \lambda_{\min} \left(v_{c,0} \mathbf{M}_{n-1}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}_{k-1}) + \sum_{j=1}^{\hat{N}} v_{c,j} \mathbf{M}_{n-1}^{\mathbf{e}_j}(\tilde{\boldsymbol{\mu}}_{k-1}) \right)$ and the corresponding normalized eigenvector \mathbf{n}_{λ} .

```
if \lambda < 0 then
```

Set $\mathcal{N}_c := \mathcal{N}_c \cup \{\mathbf{n}_\lambda\}.$

end if

end for

Solve LP (7.18) and extract the optimal solution $\tilde{\mu}_k^*$ and the optimal value \tilde{p}_k . end for

return $\langle X_i \rangle_{\rm ss}^U := \tilde{p}_{k_{\rm max}}.$
For brevity, in what follows, we will refer to LP (7.18) with $\mathcal{P} \equiv \mathcal{A}_{\binom{n+\hat{N}}{n}}$ and $\mathcal{N}_c \equiv \mathcal{A}_{\binom{n-1+\hat{N}}{n-1}}$ for all $c \in \{1, \ldots, N\}$ simply as the "initial LP".

This algorithm requires the computation of several eigenpairs consisting of a minimal eigenvalue and the corresponding eigenvector. To this end, we recommend the Hermitian Lanczos algorithm with filtering [55]. It is important to realize that the computated eigenpairs do not have to be exact. We say this for several reasons: First, only the sign of the eigenvalue is used Algorithm 4, never the actual value. Second, suppose $\mathbf{M}_n^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}^*)$ has a strictly negative eigenvalue $\lambda < 0$ with a corresponding eigenvector \mathbf{p}_{λ} . Suppose further that our computed eigenvector $\tilde{\mathbf{p}}_{\lambda}$ has some error $\Delta \mathbf{p}_{\lambda}$ such that $\tilde{\mathbf{p}}_{\lambda} = \mathbf{p}_{\lambda} + \Delta \mathbf{p}_{\lambda}$. Since $\mathbf{p}_{\lambda}^{\mathrm{T}} \mathbf{M}_n^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}) \mathbf{p}_{\lambda}$ is continuous with respect to \mathbf{p}_{λ} , the constraint $\tilde{\mathbf{p}}_{\lambda}^{\mathrm{T}} \mathbf{M}_n^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}) \tilde{\mathbf{p}}_{\lambda} \geq 0$ will still exclude $\tilde{\boldsymbol{\mu}}^*$ from the feasible set of LP (7.18) provided that the error $||\Delta \mathbf{p}_{\lambda}||$ is sufficiently small. Third, every point in the feasible set of SDP (7.12) satisfies the linear inequality $\tilde{\mathbf{p}}_{\lambda}^{\mathrm{T}} \mathbf{M}_n^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}) \tilde{\mathbf{p}}_{\lambda} \geq 0$ for any vector $\tilde{\mathbf{p}}_{\lambda}$, no matter how inaccurate. So no point will ever be "accidentally" excluded from the feasible set of LP (7.18) as a result of inaccurate eigenvalue computation.

7.4.3 Convergence

As we iterate over the values of k in Algorithm 4, each LP we solve is a *restriction* of the previous problem, in the sense that it has a smaller feasible set. Accordingly, the sequence of bounds (\tilde{p}_k) is guaranteed to be monotonically improving, in the sense that $\tilde{p}_k \leq \tilde{p}_{k-1}$, for all $k \geq 1$. It is important to note, though, that strict improvement (i.e., $\tilde{p}_k < \tilde{p}_{k-1}$) is not guaranteed in each iteration.

In the ideal scenario, the sequence of optimal values (\tilde{p}_k) generated by Algorithm 4 would converge to p^* , the optimal value of SDP (7.12). However, it is not immediately obvious that this convergence is guaranteed. In what follows, we describe a technical condition under which convergence is guaranteed. This result and its proof is based heavily on a closely related result given by Blankenship and Falk [4, Theorem 2.1].

Claim 3. If the sequence of optimal solutions $(\tilde{\mu}_k^*)$ produced by Algorithm 4 contains a convergent subsequence, then the corresponding sequence of bounds (\tilde{p}_k) converges to p^* , the optimal value of SDP (7.12).

Proof. Let $(\tilde{\boldsymbol{\mu}}_{k'}^*)$ be the convergent subsequence of $(\tilde{\boldsymbol{\mu}}_k^*)$, converging to the point $\boldsymbol{\mu}^*$. First, we aim to show that $\boldsymbol{\mu}^*$ is feasible for SDP (7.12).

By definition, each optimal solution $\tilde{\mu}_{k'}^*$ is feasible for LP (7.18), and thus satisfies $\mathbf{A}\tilde{\mu}_{k'} = \mathbf{0}$ and $\tilde{\mu}_{\mathbf{0},k'} = 1$. Since $\mathbf{A}\tilde{\mu}$ and $\tilde{\mu}_{\mathbf{0}}$ are continuous functions of $\tilde{\mu}$ it follows that

$$\lim_{k'\to\infty} \mathbf{A}\tilde{\boldsymbol{\mu}}_{k'} = \mathbf{A}\boldsymbol{\mu}^* = \mathbf{0}, \quad \text{and} \quad \lim_{k'\to\infty} = \tilde{\mu}_{\mathbf{0},k'} = \mu_{\mathbf{0}}^* = 1.$$
(7.31)

So, μ^* satisfies the linear equalities of SDP (7.12). To prove feasibility, we must show that μ^* also satisfies the LMIs. We will do so through a proof by contradiction.

Assume that μ^* violates the first LMI in the sense that $\mathbf{M}_n^0(\tilde{\mu}^*) \not\geq \mathbf{0}$ It follows that there exists some vector $\hat{\mathbf{p}} \in \mathbb{R}^{\binom{n+\hat{N}}{n}}$ such that

$$\hat{\mathbf{p}}^{\mathrm{T}}\mathbf{M}_{n}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}^{*})\hat{\mathbf{p}} < 0.$$
(7.32)

Now, the convergent subsequence of optimal solutions $(\tilde{\boldsymbol{\mu}}_{k'}^*)$ is associated with a corresponding sequence of normalized eigenvectors $(\mathbf{p}_{\lambda,k'}) \subset \mathbb{R}^{\binom{n+\hat{N}}{n}}$. Since each vector in this sequence is normalized, the $(\mathbf{p}_{\lambda,k'})$ is contained in the unit sphere – a compact set. It follows that there exists a convergent subsequence $(\mathbf{p}_{\lambda,k'_j})$ whose limit we will denote as $\hat{\mathbf{p}}_{\lambda}$. The corresponding subsequence of $(\tilde{\boldsymbol{\mu}}_{k'}^*)$ will be denoted $(\tilde{\boldsymbol{\mu}}_{k'_j}^*)$.

Since $\hat{\mathbf{p}}^{\mathrm{T}}\mathbf{M}_{n}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}})\hat{\mathbf{p}}$ is continuous with respect to $\tilde{\boldsymbol{\mu}}$, and since $\tilde{\boldsymbol{\mu}}_{k'_{j}} \to \tilde{\boldsymbol{\mu}}^{*}$, it follows from Inequality (7.32) that

$$\hat{\mathbf{p}}^{\mathrm{T}}\mathbf{M}_{n}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}_{k'j})\hat{\mathbf{p}}<0\tag{7.33}$$

for k'_i sufficiently large.

Since $\mathbf{p}_{\lambda,k'_j}$ is the normalized eigenvector corresponding to the minimum eigenvalue of $\mathbf{M}_n^0(\tilde{\boldsymbol{\mu}}_{k'_j})$, it follows that

$$\mathbf{p}_{\lambda,k_{j}'}^{\mathrm{T}}\mathbf{M}_{n}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}_{k_{j}'})\mathbf{p}_{\lambda,k_{j}'} \leq \hat{\mathbf{p}}^{\mathrm{T}}\mathbf{M}_{n}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}_{k_{j}'})\hat{\mathbf{p}}.$$
(7.34)

Taking the limit as $k'_j \to \infty$, we obtain

$$\hat{\mathbf{p}}_{\lambda}^{\mathrm{T}} \mathbf{M}_{n}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}^{*}) \hat{\mathbf{p}}_{\lambda} \leq \hat{\mathbf{p}}^{\mathrm{T}} \mathbf{M}_{n}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}^{*}) \hat{\mathbf{p}}.$$
(7.35)

With Inequality (7.32), this implies

$$\hat{\mathbf{p}}_{\lambda}^{\mathrm{T}} \mathbf{M}_{n}^{\mathbf{0}} (\tilde{\boldsymbol{\mu}}^{*}) \hat{\mathbf{p}}_{\lambda} < 0.$$
(7.36)

Now, the constraint

$$\mathbf{p}_{\lambda,k_{j}'}^{\mathrm{T}}\mathbf{M}_{n}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}})\mathbf{p}_{\lambda,k_{j}'} \geq 0$$
(7.37)

is enforced in the $(k'_j + 1)$ st LP that is solved and in every LP thereafter. Since $k'_{j+1} \geq k'_j + 1$, and since $\tilde{\mu}^*_{k'_{j+1}}$ is feasible for the (k'_{j+1}) st LP that is solved, it follows that

$$\mathbf{p}_{\lambda,k_j'}^{\mathrm{T}}\mathbf{M}_n^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}_{k_{j+1}}^*)\mathbf{p}_{\lambda,k_j'} \ge 0.$$
(7.38)

Taking the limit as $k'_j \to \infty$, we obtain

$$\hat{\mathbf{p}}_{\lambda}^{\mathrm{T}} \mathbf{M}_{n}^{\mathbf{0}} (\tilde{\boldsymbol{\mu}}^{*}) \hat{\mathbf{p}}_{\lambda} \ge 0.$$
(7.39)

This directly contradicts Inequality (7.36). Thus, we know that our original assumption, $\mathbf{M}_{n}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}^{*}) \succeq \mathbf{0}$, was false. It follows that $\mathbf{M}_{n}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}^{*}) \succeq \mathbf{0}$.

In much the same way, we could prove that $\tilde{\mu}^*$ also satisfies the other LMIs appearing in SDP (7.12).

We have shown that $\tilde{\mu}^*$ satisfies all of the constraints appearing in SDP (7.12). It follows that $\tilde{\mu}^*$ is feasible for SDP (7.12).

We now show that $\tilde{\mu}^*$ is optimal for SDP (7.12). Let \tilde{p}^* be the objective value corresponding to $\tilde{\mu}^*$. From the feasibility of $\tilde{\mu}^*$, it follows that

$$\tilde{p}^* \le p^*. \tag{7.40}$$

Since each LP solved in the execution of Algorithm 4 is a relaxation of SDP

(7.12), it follows that $\tilde{p}_{k'} \geq p^*$ for each k'. Furthermore, since $\tilde{\mu}_{k'}^* \to \tilde{\mu}^*$ and since the objective function is a continuous function of $\tilde{\mu}$, it follows that

$$\tilde{p}^* \ge p^*. \tag{7.41}$$

With Inequality (7.40), this implies $\tilde{p}^* = p^*$.

Thus, $(\tilde{\boldsymbol{\mu}}_{k'}^*)$ converges to the optimal solution of SDP (7.12), and the corresponding sequence of optimal values $(\tilde{p}_{k'})$ converges to the optimal value of SDP (7.12). Since $(\tilde{p}_{k'})$ is a subsequence of (\tilde{p}_k) and the latter sequence is monotonically decreasing, it follows that (\tilde{p}_k) converges to the optimal value of SDP (7.12) as well.

Corollary 4. If the feasible set of any LP solved in the execution of Algorithm 4 is bounded, then the sequence of bounds (\tilde{p}_k) converges to p^* , the optimal value of SDP (7.12).

Proof. Let S_k denote the feasible set of the kth LP solved in executing Algorithm 4. By construction, the feasible sets of successive LPs satisfy the relation $S_{k+1} \subset S_k$. It follows that $S_j \subset S_k$ for all $j \ge k$. Each S_k is polyhedral, and is thus closed.

Suppose that the $S_{\hat{k}}$ is bounded. Since $S_{\hat{k}}$ is also closed, it is compact. Since $S_k \subset S_{\hat{k}}$ for all $k \geq \hat{k}$, the sequence of optimal points $(\tilde{\boldsymbol{\mu}}_k^*)_{k=\hat{k}}^{\infty}$ is contained in the compact set $S_{\hat{k}}$. Every sequence contained in a compact set in Euclidean space has a convergent subsequence. In particular, $(\tilde{\boldsymbol{\mu}}_k^*)_{k=\hat{k}}^{\infty}$ has a convergent subsequence. This is still true if we prepend the finite sequence $(\tilde{\boldsymbol{\mu}}_0, \dots, \tilde{\boldsymbol{\mu}}_{k^*-1})$ to $(\tilde{\boldsymbol{\mu}}_k^*)_{k=\hat{k}}^{\infty}$, giving the full sequence $(\tilde{\boldsymbol{\mu}}_k^*)$. Then, $(\tilde{\boldsymbol{\mu}}_k^*)$ satisfies the condition of Claim 3. The conclusion follows.

7.4.4 Small Example

We now demonstrate the use of Algorithm 4 by applying it to the Michaelis-Menten system from Section 7.3.4. The results are shown in Figure 7-4. Inspecting this figure, we see that the bounds are substantially improved after just one cutting plane iteration. Applying a second iteration, they become even tighter, approaching the SDP-based bounds given in the top panel of Figure 7-1.



Figure 7-4: Bounds on the Michaelis-Menten system at steady state obtained by applying Algorithm 4. The top panel shows the bounds obtained using $k_{\text{max}} = 1$ cutting plane iterations, while the lower panel shows the bounds obtained using $k_{\text{max}} = 2$.

7.4.5 Large Example

Next, we apply Algorithm 4 by applying it to the reaction system shown in Figure 7-2, using as many as 10 cutting plane iterations. The results are shown in Figure 7-5. Comparing these bounds with those shown in Figure (7-3), it is apparent that the cutting plane iterations do, in fact, tighten the initial LP-based bounds. However, even after 10 cutting plane iterations, the gap between the LP-based bounds is still much wider than the gap between the SDP-based bounds.



Figure 7-5: Bounds on the reaction system shown in Figure 7-2 at steady state obtained by applying Algorithm 4. The top panel shows the bounds obtained using $k_{\text{max}} = 5$ cutting plane iterations, while the lower panel shows the bounds obtained using $k_{\text{max}} = 10$. Compare with the bounds shown in Figure 7-3.

7.4.6 Summary

We've seen that we can iteratively improve the bounds produced by LP (7.18), by adding vectors to the sets \mathcal{P} and \mathcal{N}_c as specified by Algorithm 4. This iterative refinement retains the numerical robustness of the original LP. Moreover, the bounds are theoretically guaranteed to monotonically improve. Unfortunately, as we saw in Section (7.4.5), the rate of improvement can be unfortunately slow.

7.5 Warm-Starting \mathcal{P} and \mathcal{N}_c

We've seen that we can improve the quality of the bounds by choosing the sets \mathcal{P} and \mathcal{N}_c appropriately. Also, while

$$\mathcal{P} \equiv \mathcal{A}_{\binom{n+\hat{N}}{n}} \tag{7.42}$$

and

$$\mathcal{N}_c \equiv \mathcal{A}_{\binom{n-1+\hat{N}}{n-1}}, \quad \forall c \in \{1, \dots, N\},$$
(7.43)

seem to be reasonable initializations of these sets based on their successful use in related literature [1], there is admittedly some arbitrariness in this choice. These two observations together suggest that it might be worth pursuing initializations of \mathcal{P} and \mathcal{N}_c which are somehow tailored for the problem at hand.

To see how this can be done, it is helpful to first understand what the LP relaxation described Section (7.3.2) means for the *dual* of SDP (7.12). In what follows, we will refer to SDP (7.12) as the *primal* SDP. We will now derive the corresponding dual SDP.

7.5.1 Derivation of the Dual SDP

This process will be made easier if we first express SDP (7.12) in terms of sums over all moments $\tilde{\mu}_{\mathbf{j}}$ that make up the vector $\tilde{\boldsymbol{\mu}}$. Let J be the set of all multi-indices $\mathbf{j} \in \mathbb{N}^{\hat{N}}$ for which there is a corresponding component $\tilde{\mu}_{\mathbf{j}}$ in $\tilde{\boldsymbol{\mu}}$. If $\mathbf{A}_{\mathbf{j}}$ represents the column of the matrix \mathbf{A} corresponding to the multi-index $\mathbf{j} \in J$, then the first constraint can be written equivalently as

$$\sum_{\mathbf{j}\in J} \mathbf{A}_{\mathbf{j}} \tilde{\mu}_{\mathbf{j}} = \mathbf{0}.$$
(7.44)

The constraint $\tilde{\mu}_{\mathbf{0}} = 1$ can be rewritten as

$$\sum_{\mathbf{j}\in J}\delta_{\mathbf{j}=\mathbf{0}}\tilde{\mu}_{\mathbf{j}} = 1.$$
(7.45)

Furthermore, the matrix $\mathbf{M}_{n}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}})$ can be written as a linear combination of coefficient matrices $\{\mathbf{M}_{n,\mathbf{j}}^{\mathbf{0}}\}_{\mathbf{j}\in J} \in \mathbb{S}^{\binom{n+\hat{N}}{n}}$:

$$\mathbf{M}_{n}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}) = \sum_{\mathbf{j}\in J} \mathbf{M}_{n,\mathbf{j}}^{\mathbf{0}} \tilde{\boldsymbol{\mu}}_{\mathbf{j}}.$$
(7.46)

So that the LMI $\mathbf{M}^{\mathbf{0}}_n(\tilde{\boldsymbol{\mu}})\succeq \mathbf{0}$ can be written equivalently as

$$\sum_{\mathbf{j}\in J} \mathbf{M}_{n,\mathbf{j}}^{\mathbf{0}} \tilde{\mu}_{\mathbf{j}} \succeq \mathbf{0}.$$
(7.47)

In much the same way, the other LMIs can be written as

$$v_{c,0}\mathbf{M}_{n-1}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}) + \sum_{j=1}^{\hat{N}} v_{c,j}\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\tilde{\boldsymbol{\mu}}) = v_{c,0}\left(\sum_{\mathbf{j}\in J}\mathbf{M}_{n-1,\mathbf{j}}^{\mathbf{0}}\tilde{\boldsymbol{\mu}}_{\mathbf{j}}\right) + \sum_{j=1}^{\hat{N}} v_{c,j}\left(\sum_{\mathbf{j}\in J}\mathbf{M}_{n-1,\mathbf{j}}^{\mathbf{e}_{j}}\tilde{\boldsymbol{\mu}}_{\mathbf{j}}\right)$$
$$= \sum_{\mathbf{j}\in J}\left(v_{c,0}\mathbf{M}_{n-1,\mathbf{j}}^{\mathbf{0}} + \sum_{j=1}^{\hat{N}} v_{c,j}\mathbf{M}_{n-1,\mathbf{j}}^{\mathbf{e}_{j}}\right)\tilde{\boldsymbol{\mu}}_{\mathbf{j}} \succeq \mathbf{0}, \quad \forall c \in \{1,\ldots,N\}.$$
(7.48)

Finally, the objective function can be rewritten as

$$v_{i,0} + \sum_{j=1}^{\hat{N}} v_{i,j} \tilde{\mu}_{\mathbf{e}_j} = \sum_{\mathbf{j} \in J} \left(v_{i,0} \delta_{\mathbf{j}=\mathbf{0}} + \sum_{j=1}^{\hat{N}} v_{i,j} \delta_{\mathbf{j}=\mathbf{e}_j} \right) \tilde{\mu}_{\mathbf{j}}.$$
 (7.49)

Substituting these expressions into SDP (7.12), we obtain

$$p^{*} = \max_{\tilde{\boldsymbol{\mu}}} \sum_{\mathbf{j} \in J} \left(v_{i,0} \delta_{\mathbf{j}=\mathbf{0}} + \sum_{j=1}^{\tilde{N}} v_{i,j} \delta_{\mathbf{j}=\mathbf{e}_{j}} \right) \tilde{\mu}_{\mathbf{j}}$$

s.t.
$$\sum_{\mathbf{j} \in J} \mathbf{A}_{\mathbf{j}} \tilde{\mu}_{\mathbf{j}} = \mathbf{0},$$
$$\sum_{\mathbf{j} \in J} \delta_{\mathbf{j}=\mathbf{0}} \tilde{\mu}_{\mathbf{j}} = 1,$$
$$\sum_{\mathbf{j} \in J} \mathbf{M}_{n,\mathbf{j}}^{\mathbf{0}} \tilde{\mu}_{\mathbf{j}} \succeq \mathbf{0},$$
$$\sum_{\mathbf{j} \in J} \mathbf{M}_{n,\mathbf{j}}^{\mathbf{0}} \tilde{\mu}_{\mathbf{j}} \succeq \mathbf{0},$$
$$\sum_{\mathbf{j} \in J} \left(v_{c,0} \mathbf{M}_{n-1,\mathbf{j}}^{\mathbf{0}} + \sum_{j=1}^{\tilde{N}} v_{c,j} \mathbf{M}_{n-1,\mathbf{j}}^{\mathbf{e}_{j}} \right) \tilde{\mu}_{\mathbf{j}} \succeq \mathbf{0}, \quad \forall c \in \{1, \dots, N\}.$$

The next step is to introduce a dual multiplier for each constraint, and then bring

that constraint into the objective function. Doing so, the problem becomes

$$U(\mathbf{q}, u, \mathbf{P}, (\mathbf{N}_{c})_{c=1}^{N}) \equiv \max_{\tilde{\boldsymbol{\mu}}} \sum_{\mathbf{j} \in J} \left(v_{i,0} \delta_{\mathbf{j}=\mathbf{0}} + \sum_{j=1}^{\hat{N}} v_{i,j} \delta_{\mathbf{j}=\mathbf{e}_{j}} \right) \tilde{\mu}_{\mathbf{j}} + \mathbf{q}^{\mathrm{T}} \left(\sum_{\mathbf{j} \in J} \mathbf{A}_{\mathbf{j}} \tilde{\mu}_{\mathbf{j}} \right) \\ + u \left(1 - \sum_{\mathbf{j} \in J} \delta_{\mathbf{j}=\mathbf{0}} \tilde{\mu}_{\mathbf{j}} \right) \\ + \mathbf{P} \cdot \left(\sum_{\mathbf{j} \in J} \mathbf{M}_{n,\mathbf{j}}^{\mathbf{0}} \tilde{\mu}_{\mathbf{j}} \right) \\ + \sum_{c=1}^{N} \mathbf{N}_{c} \cdot \left(\sum_{\mathbf{j} \in J} \left(v_{c,0} \mathbf{M}_{n-1,\mathbf{j}}^{\mathbf{0}} + \sum_{j=1}^{\hat{N}} v_{c,j} \mathbf{M}_{n-1,\mathbf{j}}^{\mathbf{e}_{j}} \right) \tilde{\mu}_{\mathbf{j}} \right),$$
(7.51)

where the "dot" notation is the standard matrix inner product for symmetric matrices, defined by $\mathbf{A} \cdot \mathbf{B} \equiv \operatorname{tr}(\mathbf{AB})$. If the matrices \mathbf{P} and $(\mathbf{N}_c)_{c=1}^N$ are positive semidefinite, then

$$U(\mathbf{q}, u, \mathbf{P}, (\mathbf{N}_c)_{c=1}^N) \ge p^*, \tag{7.52}$$

by the standard duality argument. Specifically, any point $\tilde{\mu}$ which is feasible for SDP (7.50) is also feasible for Problem (7.51). Moreover, if $\tilde{\mu}$ which is feasible for SDP (7.50), it follows that

$$\sum_{\mathbf{j}\in J} \mathbf{A}_{\mathbf{j}} \tilde{\mu}_{\mathbf{j}} = \mathbf{0},\tag{7.53}$$

$$1 - \sum_{\mathbf{j} \in J} \delta_{\mathbf{j}=\mathbf{0}} \tilde{\mu}_{\mathbf{j}} = 0, \qquad (7.54)$$

$$\sum_{\mathbf{j}\in J} \mathbf{M}_{n,\mathbf{j}}^{\mathbf{0}} \tilde{\mu}_{\mathbf{j}} \succeq \mathbf{0}, \tag{7.55}$$

$$\sum_{\mathbf{j}\in J} \left(v_{c,0} \mathbf{M}_{n-1,\mathbf{j}}^{\mathbf{0}} + \sum_{j=1}^{\hat{N}} v_{c,j} \mathbf{M}_{n-1,\mathbf{j}}^{\mathbf{e}_j} \right) \tilde{\mu}_{\mathbf{j}} \succeq \mathbf{0}.$$
(7.56)

From LMI (7.55) and $\mathbf{P} \succeq \mathbf{0}$, we are guaranteed that

$$\mathbf{P} \cdot \left(\sum_{\mathbf{j} \in J} \mathbf{M}_{n,\mathbf{j}}^{\mathbf{0}} \tilde{\mu}_{\mathbf{j}} \right) \ge 0.$$
(7.57)

Similarly, from LMI (7.56) and $\mathbf{N}_c \succeq \mathbf{0}$ for all $c \in \{1, \ldots, N\}$, we are guaranteed that

$$\sum_{c=1}^{N} \mathbf{N}_{c} \cdot \left(\sum_{\mathbf{j} \in J} \left(v_{c,0} \mathbf{M}_{n-1,\mathbf{j}}^{\mathbf{0}} + \sum_{j=1}^{\hat{N}} v_{c,j} \mathbf{M}_{n-1,\mathbf{j}}^{\mathbf{e}_{j}} \right) \tilde{\mu}_{\mathbf{j}} \right) \ge 0.$$
(7.58)

It follows that the objective value associated with $\tilde{\mu}$ in Problem (7.51) is greater than or equal to the objective value associated with $\tilde{\mu}$ in SDP (7.50). Since this is true for every $\tilde{\mu}$ which is feasible for SDP (7.50), Inequality (7.52) follows.

Algebraically manipulating the objective function, Problem (7.51) can be rewritten as

$$\max_{\hat{\mu}} \sum_{\mathbf{j}\in J} \left(v_{i,0}\delta_{\mathbf{j}=\mathbf{0}} + \sum_{j=1}^{\hat{N}} v_{i,j}\delta_{\mathbf{j}=\mathbf{e}_{j}} \right) \tilde{\mu}_{\mathbf{j}} \\
+ \sum_{\mathbf{j}\in J} \left(\mathbf{q}^{\mathrm{T}}\mathbf{A}_{\mathbf{j}} \right) \tilde{\mu}_{\mathbf{j}} \\
+ u - \sum_{\mathbf{j}\in J} \left(u\delta_{\mathbf{j}=\mathbf{0}} \right) \tilde{\mu}_{\mathbf{j}} \\
+ \sum_{\mathbf{j}\in J} \left(\mathbf{P} \cdot \mathbf{M}_{n,\mathbf{j}}^{\mathbf{0}} \right) \tilde{\mu}_{\mathbf{j}} \\
+ \sum_{\mathbf{j}\in J} \left(\sum_{c=1}^{N} \mathbf{N}_{c} \cdot \left(v_{c,0}\mathbf{M}_{n-1,\mathbf{j}}^{\mathbf{0}} + \sum_{j=1}^{\hat{N}} v_{c,j}\mathbf{M}_{n-1,\mathbf{j}}^{\mathbf{e}_{j}} \right) \right) \tilde{\mu}_{\mathbf{j}}$$

$$(7.59)$$

$$= \max_{\hat{\mu}} u + \sum_{\mathbf{j}\in J} \begin{pmatrix} v_{i,0}\delta_{\mathbf{j}=\mathbf{0}} + \sum_{j=1}^{\hat{N}} v_{i,j}\delta_{\mathbf{j}=\mathbf{e}_{j}} \\
+ \mathbf{q}^{\mathrm{T}}\mathbf{A}_{\mathbf{j}} \\
- u\delta_{\mathbf{j}=\mathbf{0}} \\
+ \mathbf{P} \cdot \mathbf{M}_{n,\mathbf{j}}^{\mathbf{0}} \\
+ \sum_{c=1}^{N} \mathbf{N}_{c} \cdot \left(v_{c,0}\mathbf{M}_{n-1,\mathbf{j}}^{\mathbf{0}} + \sum_{j=1}^{\hat{N}} v_{c,j}\mathbf{M}_{n-1,\mathbf{j}}^{\mathbf{e}_{j}} \right) \\ \tilde{\mu}_{\mathbf{j}}.$$

Let the term in the large parentheses be represented concisely as $w_j(\mathbf{q}, u, \mathbf{P}, (\mathbf{N}_c)_{c=1}^N)$

for each $\mathbf{j} \in J$. Then, Problem (7.51) can be written relatively concisely as

$$U(\mathbf{q}, u, \mathbf{P}, (\mathbf{N}_c)_{c=1}^N) = \max_{\tilde{\boldsymbol{\mu}}} \left\{ u + \sum_{\mathbf{j} \in J} w_{\mathbf{j}}(\mathbf{q}, u, \mathbf{P}, (\mathbf{N}_c)_{c=1}^N) \tilde{\boldsymbol{\mu}}_{\mathbf{j}} \right\}.$$
 (7.60)

Since $\tilde{\mu}$ is a free variable in the optimization problem, this last expression implies that

$$U(\mathbf{q}, u, \mathbf{P}, (\mathbf{N}_c)_{c=1}^N) = \begin{cases} u & \text{if } w_{\mathbf{j}}(\mathbf{q}, u, \mathbf{P}, (\mathbf{N}_c)_{c=1}^N) = 0, \quad \forall \mathbf{j} \in J, \\ +\infty & \text{otherwise.} \end{cases}$$
(7.61)

Collecting our results, we've seen

1. If the matrices \mathbf{P} and $(\mathbf{N}_c)_{c=1}^N$ are positive semidefinite, then

$$U(\mathbf{q}, u, \mathbf{P}, (\mathbf{N}_c)_{c=1}^N) \ge p^*.$$
 (7.62)

2. If $w_{\mathbf{j}}(\mathbf{q}, u, \mathbf{P}, (\mathbf{N}_c)_{c=1}^N) = 0$ for all $\mathbf{j} \in J$, then

$$U(\mathbf{q}, u, \mathbf{P}, (\mathbf{N}_c)_{c=1}^N) = u.$$
(7.63)

It follows that if we can find some $\mathbf{q}, u, \mathbf{P}$ and $(\mathbf{N}_c)_{c=1}^N$ such that

$$w_{\mathbf{j}}(\mathbf{q}, u, \mathbf{P}, (\mathbf{N}_c)_{c=1}^N) = 0, \quad \forall \mathbf{j} \in J,$$
(7.64)

$$\mathbf{P} \succeq \mathbf{0},\tag{7.65}$$

and

$$\mathbf{N}_c \succeq \mathbf{0}, \quad \forall c \in \{1, \dots, N\},\tag{7.66}$$

then $u \ge p^*$ (i.e. u is an upper bound on p^*). Ideally, we want this upper bound to

be as tight as possible, which leads us to the the following optimization problem:

$$d^* \equiv \min_{\substack{\mathbf{q}, u, \mathbf{P}, \\ (\mathbf{N}_c)_{c=1}^N}} u$$

s.t. $w_{\mathbf{j}}(\mathbf{q}, u, \mathbf{P}, (\mathbf{N}_c)_{c=1}^N) = 0, \quad \forall \mathbf{j} \in J,$
 $\mathbf{P} \succeq \mathbf{0},$
 $\mathbf{N}_c \succeq \mathbf{0}, \quad \forall c \in \{1, \dots, N\}.$ (7.67)

Recalling that the expression $w_{\mathbf{j}}(\cdot, \cdot, \cdot, \cdot)$ is affine with respect to each of its arguments, we see that this is an SDP. More specifically, it is the dual of SDP (7.12).

By construction, we are guaranteed that $d^* \ge p^*$. Our computational experience suggests that there is no duality gap (i.e., that $d^* = p^*$). However, we have not yet been able to prove this. In any case, since p^* is an upper bound on $\langle X_i \rangle_{ss}$, the steady-state mean molecular count of species $i \in \{1, \ldots, N\}$, it follows that d^* is also an upper bound on $\langle X_i \rangle_{ss}$.

7.5.2 Sum of Squared Polynomials Interpretation

The dual SDP (7.67) can be interpreted as a *polynomial optimization problem*, where the constraints enforce the nonegativity of a polynomial on the set $\bar{\mathcal{X}}$. This interpretation is not essential for understanding what follows. We note it now simply because it is interesting, and future work may benefit from this interpretation.

7.5.3 Dual Equivalent of the LP Relaxation of the Primal SDP

One could repeat the derivation of Section 7.5.1, but, instead of starting from SDP (7.12), start from its LP relaxation, (7.18). The resulting dual LP is

$$\begin{split} \tilde{d} &\equiv \min_{\substack{\mathbf{q}, u, \mathbf{P}, \\ (\mathbf{N}_c)_{c=1}^N, \\ \alpha, \beta}} u \\ \text{s.t.} \quad w_{\mathbf{j}}(\mathbf{q}, u, \mathbf{P}, (\mathbf{N}_c)_{c=1}^N) = 0, \quad \forall \mathbf{j} \in J, \\ \mathbf{P} &= \sum_{\mathbf{p} \in \mathcal{P}} \alpha_{\mathbf{p}} \mathbf{p} \mathbf{p}^{\mathrm{T}}, \\ \mathbf{N}_c &= \sum_{\mathbf{n} \in \mathcal{N}_c} \beta_{c, \mathbf{n}} \mathbf{n} \mathbf{n}^{\mathrm{T}}, \quad \forall c \in \{1, \dots, N\}, \\ \alpha_{\mathbf{p}} \geq 0, \quad \forall \mathbf{p} \in \mathcal{P}, \\ \beta_{c, \mathbf{n}} \geq 0, \quad \forall \mathbf{n} \in \mathcal{N}_c, \quad \forall c \in \{1, \dots, N\}. \end{split}$$
(7.68)

An Inner Approximation

Comparing the dual LP (7.68) with the dual SDP (7.67), we see that the two are very similar, except that dual LP (7.68) approximates each positive semidefinite matrix appearing in dual SDP (7.67) with a conic combination of finitely many rank one, positive semidefinite matrices. Of course, generally speaking, any positive semidefinite matrix $\mathbf{Y} \in \mathbb{S}^k$ can be written as

$$\mathbf{Y} = \sum_{i=1}^{k} \lambda_i \mathbf{y}_i \mathbf{y}_i^{\mathrm{T}}, \quad \lambda_i \ge 0, \quad \forall i \in \{1, \dots, k\},$$
(7.69)

for some set of vectors $\{\mathbf{y}_i\}_{i=1}^k \equiv \mathcal{Y}$. However, remembering that \mathcal{P} and $(\mathcal{N}_c)_{c=1}^N$ are fixed, finite sets of prespecified vectors, we see that dual LP (7.68) is a *restriction* of dual SDP (7.67). From any feasible point of dual LP (7.68) it is possible to construct a feasible point for dual SDP (7.67). However, it is not necessarily possible to go the other way – using a feasible point for dual SDP (7.67) to construct a feasible point for dual LP (7.68). Thus, in some sense, the feasible set of dual LP (7.68) is an *inner* approximation of the feasible set of dual SDP (7.67). Given that both problems share the same objective value, this implies that $\tilde{d} \ge d^*$. Since $d^* \ge \langle X_i \rangle_{ss}$, it follows that $\tilde{d} \ge \langle X_i \rangle_{ss}$.

Comparison with the Primal LP

We wish to remind the reader that while dual LP (7.68) is a restriction of dual SDP (7.67), in the primal picture, LP (7.18) is a relaxation of SDP (7.12).

By strong duality, we know that the optimal values of primal LP (7.18) and dual LP (7.68) are equal (i.e., that $\tilde{p} = \tilde{d}$). Thus, in some sense, the two problems are equivalent.

Notice that for every inequality constraint appearing in the primal LP (7.18) we have a corresponding rank one matrix and nonnegative decision variable in the dual LP (7.68). Thus, every time we add a cutting plane to primal LP (7.18) in Algorithm 4, we are also adding a decision variable to dual LP (7.68). This observation is in agreement with the general fact that, for linear programs, adding a cutting plane to the primal problem is equivalent to "column generation" for the dual problem [1].

The two problems are also complementary. While the primal perspective can be used to motivate the cutting plane procedure, the dual perspective can guide us towards a set of ideal cutting planes. We will explore this idea further in the next section.

The Perfect Choice of \mathcal{P} and \mathcal{N}_c

Assume that dual SDP (7.67) has an optimal solution $(\mathbf{q}^*, u^*, \mathbf{P}^*, (\mathbf{N}_c^*)_{c=1}^N)$, with a corresponding optimal value of $d^* = u^*$. By definition, \mathbf{P}^* must satisfy $\mathbf{P}^* \succeq \mathbf{0}$. It follows that \mathbf{P}^* can be written as

$$\mathbf{P}^* = \sum_{k=1}^{\binom{n+N}{n}} \alpha_k^* \mathbf{p}_k \mathbf{p}_k^{\mathrm{T}}, \qquad (7.70)$$

for some set of vectors $\{\mathbf{p}_k\}_{k=1}^{\binom{n+\hat{N}}{n}}$ and a corresponding set of nonnegative coefficients $\alpha_k^* \geq 0$. Similarly, for each $c \in \{1, \ldots, N\}$, we have

$$\mathbf{N}_{c}^{*} = \sum_{k=1}^{\binom{n-1+\hat{N}}{n-1}} \beta_{c,k}^{*} \mathbf{n}_{c,k} \mathbf{n}_{c,k}^{\mathrm{T}}, \qquad (7.71)$$

for some set of vectors $\{\mathbf{n}_{c,k}\}_{k=1}^{\binom{n-1+\hat{N}}{n-1}}$ and a corresponding set of nonnegative coefficients $\beta_{c,k}^* \geq 0$.

Suppose that we set

$$\mathcal{P} \equiv \{\mathbf{p}_k\}_{k=1}^{\binom{n+N}{n}} \tag{7.72}$$

and

$$\mathcal{N}_c \equiv \{\mathbf{n}_{c,k}\}_{k=1}^{\binom{n-1+\hat{N}}{n-1}} \tag{7.73}$$

and consider dual LP (7.68). Setting $\mathbf{q} \equiv \mathbf{q}^*$, $u \equiv u^*$, $\mathbf{P} = \mathbf{P}^*$, and $\mathbf{N}_c = \mathbf{N}_c^*$ for each $c \in \{1, \ldots, N\}$, we are guaranteed that we satisfy the first family of equality constraints:

$$w_{\mathbf{j}}(\mathbf{q}, u, \mathbf{P}, (\mathbf{N}_c)_{c=1}^N) = 0, \quad \forall \mathbf{j} \in J.$$
(7.74)

Furthermore, setting

$$\alpha_{\mathbf{p}_k} \equiv \alpha_k^*, \quad \forall k \in \left\{1, \dots, \binom{n+\hat{N}}{n}\right\}$$
(7.75)

and

$$\beta_{c,\mathbf{n}_{c,k}} \equiv \beta_{c,k}^*, \quad \forall k \in \left\{1, \dots, \binom{n-1+\hat{N}}{n-1}\right\}, \ \forall c \in \{1, \dots, N\},$$
(7.76)

we are guaranteed that we satisfy the remaining constraints. We have thus constructed a feasible point for dual LP (7.68). Furthermore, the objective value associated with this feasible point is $u^* = d^*$. Since $\tilde{d} \ge d^*$, it follows that the feasible point we have constructed is not only feasible but also optimal for LP (7.68).

Thus, for this particular choice of \mathcal{P} and $\{\mathcal{N}_c\}_{c=1}^N$, the optimal value of dual LP

(7.68) equals the optimal value of dual SDP (7.67). Given that dual LP (7.68) is a restriction of dual SDP (7.67), this is the best possible outcome. Thus, in some sense, the choice of \mathcal{P} and $\{\mathcal{N}_c\}_{c=1}^N$ that we have described above is perfect.

Notice, also, that if we were to augment the perfect sets \mathcal{P} and $\{\mathcal{N}_c\}_{c=1}^N$ with extraneous vectors, setting the corresponding coefficients equal to zero, we would still obtain a feasible point with an optimal value equal to d^* .

Since primal LP (7.18) and dual LP (7.68) are equivalent (i.e., $\tilde{p} = \tilde{d}$), it follows from the above discussion that the perfect choice of \mathcal{P} and $\{\mathcal{N}_c\}_{c=1}^N$ for dual LP (7.68) is also the perfect choice for primal LP (7.18). Thus, the ideal cutting planes for primal LP (7.18) can be obtained from the optimal solution of dual SDP (7.67).

An Imperfect Choice of \mathcal{P} and \mathcal{N}_c

Suppose that instead of the perfect vectors $\{\mathbf{p}_k\}_{k=1}^{\binom{n+\hat{N}}{n}}$ and $\{\mathbf{n}_{c,k}\}_{k=1}^{\binom{n-1+\hat{N}}{n-1}}$, $\forall c \in \{1, \ldots, N\}$, we have a collection of slightly perturbed vectors

$$\left\{\tilde{\mathbf{p}}_k\right\}_{k=1}^{\binom{n+\hat{N}}{n}} \tag{7.77}$$

and

$$\{\tilde{\mathbf{n}}_{c,k}\}_{k=1}^{\binom{n-1+\hat{N}}{n-1}}, \quad \forall c \in \{1, \dots, N\}$$
(7.78)

with which we define our sets \mathcal{P} and $\{\mathcal{N}_c\}_{c=1}^N$. Then, we are not guaranteed that primal LP (7.18) (or dual LP (7.68)) achieves an optimal value of d^* . However, assuming that the perturbation from perfection is not too large, we can still reasonably expect the perturbed vectors to provide decent cutting planes in primal LP (7.18). To justify this statement, we again appeal to the fact that all of the cutting planes are based on quadratic forms, such as $\mathbf{p}^{\mathrm{T}}\mathbf{M}_{n}^{0}(\tilde{\boldsymbol{\mu}})\mathbf{p}$, and that this quadratic form is a continuous function of \mathbf{p} .

Technically, attempting to solve primal LP (7.18) with the perturbed vectors described above, we are not guaranteed that the problem has a bounded solution. However, as explained in Section 7.3.3, for stochastic chemical kinetic systems where the set $\bar{\mathcal{X}}$ is bounded, inclusion of the Ahmadi vectors in the sets \mathcal{P} and $\{\mathcal{N}_c\}_{c=1}^N$ is

sufficient to guarantee that primal LP (7.18) is bounded¹. Thus, it is reasonable to set

$$\mathcal{P} \equiv \{\tilde{\mathbf{p}}_k\}_{k=1}^{\binom{n+\hat{N}}{n}} \cup \mathcal{A}_{\binom{n+\hat{N}}{n}},\tag{7.79}$$

and

$$\mathcal{N}_{c} \equiv \{ \tilde{\mathbf{n}}_{c,k} \}_{k=1}^{\binom{n-1+\hat{N}}{n-1}} \cup \mathcal{A}_{\binom{n-1+\hat{N}}{n-1}}, \quad \forall c \in \{1, \dots, N\}.$$
(7.80)

7.5.4 The Warm-Starting Algorithm

Suppose that we attempt to solve primal SDP (7.12) with an SDP solver such as SeDuMi, and we obtain a solution with a certificate of optimality. In this case, we are done, we can take the optimal value p^* as an upper bound on $\langle X_i \rangle_{ss}$.

On the other hand, suppose that the computed solution is inaccurate in that it does not satisfy the solver's numerical tolerances. Then, the computed optimal value cannot be trusted as an upper bound on $\langle X_i \rangle_{ss}$. However, the effort put into obtaining this inaccurate solution need not be wasted. Along with the inaccurate primal optimal solution, $\tilde{\mu}^*$, any SDP solver using a state-of-the-art primal-dual interior point method will also return an approximate dual optimal solution $(\mathbf{q}^*, u^*, \mathbf{P}^*, (\mathbf{N}_c^*)_{c=1}^N)$. We can perform eigenvalue decompositions for the matrices \mathbf{P}^* and $\{\mathbf{N}_c^*\}_{c=1}^N$ to obtain approximations of the vectors that constitute the perfect choice of \mathcal{P} and $\{\mathcal{N}_c\}_{c=1}^N$. Supplementing these vectors with the Ahmadi vectors, we can then solve primal LP (7.18), expecting to obtain reasonably good bounds. Of course, the bounds may not be quite as tight as desired. In this case, we have the option of applying cutting plane iterations to refine the already good (but not perfect) bounds.

What we have just described can be viewed as a method for warm-starting Algorithm 4 with sets \mathcal{P} and $\{\mathcal{N}_c\}_{c=1}^N$ that are tailored to the problem at hand. Alternatively, it can be viewed as using an SDP solver to obtain a quick but inaccurate solution to primal SDP (7.12), and then using a more numerically robust LP solver to find a "nearby" solution which provides a guaranteed bound.

¹Actually, from the argument in Section 7.3.3, only the inclusion of the vector \mathbf{e}_1 in each set \mathcal{N}_c is necessary. However, we include the rest of the vectors "for good measure". They can only improve the quality of the bound.

These ideas are formalized in Algorithm 5.

Algorithm 5 Sequential SDP-LP Solve Input:

- 1. The index $i \in \{1, ..., N\}$ corresponding to the species whose mean you'd like to bound.
- 2. The matrix **A**.
- 3. The vectors $\{\mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_{\hat{N}}\}$
- 4. The maximum number of cutting plane iterations, k_{max} .

Output: An upper bound $\langle X_i \rangle_{ss}^U$ on the steady-state mean molecular count of species *i*.

Algorithm:

Attempt to solve primal SDP (7.12).

if the solution is numerically accurate then

return $\langle X_i \rangle_{ss}^U := p^*$.

else

Extract the matrices \mathbf{P}^* and $\{\mathbf{N}_c^*\}_{c=1}^N$ from the inaccurate dual optimal solution.

Set
$$\mathcal{P} := \mathcal{A}_{\binom{n+\hat{N}}{n}}$$
 and $\mathcal{N}_c := \mathcal{A}_{\binom{n-1+\hat{N}}{n-1}}$ for all $c \in \{1, \dots, N\}$.

Augment \mathcal{P} with the eigenvectors of \mathbf{P}^* .

Augment \mathcal{N}_c with the eigenvectors of \mathbf{N}_c^* , for all $c \in \{1, \ldots, N\}$.

Solve LP (7.18) and extract the optimal solution $\tilde{\mu}^*$ and the optimal value \tilde{p} .

Apply cutting plane iterations, as described in Algorithm 4.

return $\langle X_i \rangle_{ss}^U := \tilde{p}.$

end if

7.5.5 Example

To demonstrate the utility of the sequential SDP-LP solve, we now apply it to a reaction system representative of a gene network with autoregulatory negative feedback:

$$G \xrightarrow{c_{1}} G + M$$

$$M \xrightarrow{c_{2}} P + M$$

$$P + E \xleftarrow{c_{3}} EP$$

$$EP \xrightarrow{c_{5}} E$$

$$M \xrightarrow{c_{6}} \emptyset$$

$$P + G \xleftarrow{c_{7}} GP$$

$$P + GP \xleftarrow{c_{9}} GP_{2}$$

$$GP \xrightarrow{c_{1}} GP + M$$

$$(7.81)$$

This system was first proposed and studied by Thomas et al. [68]. Species M is intended to represent messenger RNA (mRNA), produced by transcription from a single gene, G. The mRNA is translated into a protein, P, which is subsequently degraded by enzyme E. The protein can also reversibly bind to the gene at two sites. If only one site is occupied, giving the species GP, the mRNA continues to be transcribed. If both sites are occupied, giving species GP_2 , no transcription can take place.

Reaction System (7.81) has 7 distinct chemical species and 11 reactions. In this respect, it is smaller than the reaction system described in Section 7.3.5. However, while the previous reaction system had finitely many reachable states, Reaction System (7.81) has infinitely many reachable states.

Thomas et al. [68] specify the rate constants, k_r , for reactions 2 through 10, which can be translated into the microscropic rate constants, c_r , in Table 7.1. We assume an initial condition of $x_{\rm G} = 1$, $x_{\rm E} = 100$, and all other species zero, which is consistent with the total molecular counts for the gene and enzyme specified by Thomas et al. [68].

| Reaction Index, r | $c_r ({\rm s}^{-1})$ |
|---------------------|-----------------------|
| 2 | 1×10^{-2} |
| 3 | 1×10^{-6} |
| 4 | 1×10^{-1} |
| 5 | 1×10^{-2} |
| 6 | 1×10^{-2} |
| 7 | 1×10^{-9} |
| 8 | 1×10^{-1} |
| 9 | 1×10^{-2} |
| 10 | 1×10^{-1} |
| | |

Table 7.1: Microscopic rate constants for the autoregulatory gene network.

As with the previous examples, we conducted a parameter sweep, varying c_1 from 1×10^{-4} to 1×10^{-3} . For each value of c_1 , we attempted to calculate bounds on $\langle P \rangle_{ss}$ by solving SDP (7.12). For several of the bounds, CVX returned an "inaccurate" warning flag, casting doubt on the validity of the results. We then applied Algorithm 4, with $k_{max} = 5$. Each LP was solved successfully. However, there was a wide gap in the bounds for several values of c_1 , as shown in the top panel of Figure 7-6. In contrast, applying Algorithm 5 gave the very tight bounds shown in the bottom panel of Figure 7-6. Notably, no cutting planes were required to generate these bounds. Only one LP was solved after each inaccurate SDP solve. Each of these LPs was solved accurately, to Gurobi's default precision.



Figure 7-6: Bounds on $\langle P \rangle_{ss}$ generated through Algorithm 4 (top panel) and Algorithm 5 (bottom panel).

7.6 Conclusion

In this chapter, we have described an algorithm for improving the numerical reliability of the bounding method. This algorithm is based on LPs which outer- and inner-approximate the original primal and dual SDPs, respectively. We've shown that the approximating LPs can be iteratively improved through the computation of cutting planes, and intelligently initialized using the results of an inaccurate SDP solve. Finally, we have demonstrated the efficacy of the algorithm by applying it to a stochastic chemical kinetics system taken from the literature which models an autoregulatory gene network.

While we have discussed the algorithm specifically in the context of the problem of calculating bounds on the steady-state mean molecular count of a species, it could readily be generalized to bounds on other quantities and to the dynamic analysis.

Chapter 8

Conclusion

8.1 Summary

This thesis demonstrates how, given a partial description of a distribution, we can use SDPs to calculate mathematically rigorous bounds on other descriptions. We considered two distinct problems, briefly considering particle size distributions and more extensively analyzing the problem of stochastic chemical kinetics. We demonstrated that moment-based SDPs provide an effective means of dealing with the closure problem as it appears in stochastic chemical kinetics. We showed that the SDP bounding method, developed in the context of the steady-state problem, could be generalized to the dynamic problem. Furthermore, we showed that our bounding method can be used to analyze stochastic chemical kinetic systems with features that frustrate analysis through other methods. Several refinements of the bounding method were presented which aim to mitigate either theoretical or numerical weaknesses of the original method.

8.2 Future Directions

As noted throughout the preceding chapters, there are still many lingering questions, requiring further research. We now note a few other research areas which may prove to be fruitful. While we focused on solving the moment closure problem in the context of stochastic chemical kinetics, it also presents itself in population balance models describing other systems (e.g., aggregrating particles). It would be interesting to apply the SDP bounding method to these systems as well.

In our formulation of the dynamic bounding SDP, we had a distinct set of LMIs for each value of $\rho \in \mathcal{R}$. The corresponding vectors $\mathbf{z}^{(\rho)}$ were related only indirectly through $\mathbf{y}(T)$. However, if we were to pick our values of ρ so that they were equally spaced in \mathbb{R}_- , it would be possible to construct a set of LMIs involving all values of $\rho \in \mathcal{R}$ simultaneously. This would actually be a closer analogy of what is done in the Bertsimas and Caramanis paper [3] than what we have currently proposed. Constructing these larger LMIs would create more direct dependencies between the $\mathbf{z}^{(\rho)}$ vectors, further restricting them, and potentially producing tighter bounds. This suggestion contradicts our previous suggestion that the values of \mathcal{R} should be picked to approximate eigenvalues of the matrix \mathbf{G} . However, maybe some compromise is possible, in which for each approximated eigenvalue $\rho \in \mathcal{R}$, we also have, for example, $\{\frac{1}{5}\rho, \frac{2}{5}\rho, \frac{3}{5}\rho, \frac{4}{5}\rho\} \subset \mathcal{R}$.

One of the unappealing features of the extents-based reformulation of the bounding SDP described in Chapter 5 is that it introduces redundancy in the sense that there are multiple states $\boldsymbol{\epsilon} \in \mathcal{E}$ corresponding to a single state $\mathbf{x} \in \mathcal{X}$. This likely leads to degeneracy in the resulting SDPs, which exacerbates an already serious numerical problem. To eliminate this redundancy, we propose that some reactions be eliminated and others designated as reversible (with ϵ_r values allowed to be negative) so that each $\mathbf{x} \in \mathcal{X}$ corresponds to a unique $\boldsymbol{\epsilon} \in \mathcal{E}$. Our initial thoughts on this suggest that it could be accomplished with some modified form of Gaussian Elimination.

While the LP approximations described in Chapter 6 do improve the numerical tractability of the bounding method, so that it can be successfully applied to larger more complex systems, numerical limitations are still a concern. We expect that some of these numerical issues will be alleviated over time, as SDP solver algorithms and codes become more sophisticated and robust. However, it seems that our prototypical bounding SDP is intrinsically ill-conditioned, in which case, improvement in the SDP

solvers will not be enough. It may be that this ill-conditioning is only alleviated by some clever reformulation of the bounding SDP, which, up until now, has alluded us. Our best guess is that the answer lies in the replacing the moments based on momonials with moments based on some set of orthogonal polynomials. Of course, orthogonal polynomials are always defined with respect to some distribution, and it is unclear a priori which distribution and thus which orthogonal polynomials would be the best choice. Perhaps the most appropriate distribution is that which maximizes or minimizes the quantity of interest. In that case, the answer may involve an iterative scheme, in which the result of an approximate SDP solve is used generate the set of orthogonal polynomials. This is currently nothing more than speculation, but it may be worth further exploration.

Appendix A

Supplemental Material for Chapter 2

This is a recreation of the appendix published with the paper[12] on which Chapter 2 is based.

A.0.1 Implementation Details

All numerical examples in this paper were computed on a 64-bit Dell Precision T3610 workstation with a 3.70 GHz Intel Xeon CPU. In each example, CVX [28] was used to model the SDP, using the default tolerance (i.e. "precision") settings. SeDuMi [66] was used as the underlying solver.

A.0.2 The Problem with LP-based Methods

In the introduction, we noted that LP-based methods for calculating bounds on PSDs are "not truly rigorous". By this, we mean that the bounds they produce are not necessarily valid. There may exist a distribution which satisfies all the known data (i.e., constraints) and yet violates the supposed bound.

The reason for this is that the LP-based methods optimize only over those distributions supported on finitely many predetermined "grid locations", r_1, \ldots, r_n . This means that we are only considering distributions with particle sizes r_1, \ldots, r_n . No other sizes are permitted. However, it is likely that the distribution which maximizes/minimizes the quantity we care about includes particles of size s outside of the predetermined set $\{r_1, \ldots, r_n\}$. In this case, the bound computed by considering only distributions on r_1, \ldots, r_n is invalid. It is a false bound.

We will illustrate this point by applying the LP-based method to the problem appearing in Example 2.3.1. In that example, we wished to compute an upper bound on the number of particles with size in the range 85 μ m to 150 μ m, with the moments given in (2.19). Following the paradigm described in [45], we formulate the following LP:

$$N_{\text{LP}}^{U} \equiv \max_{\mathbf{w}} \quad \mathbf{c}^{\text{T}} \mathbf{w}$$

s.t.
$$\sum_{i=1}^{n} w_{i} r_{i}^{k} = \mu_{k}, \text{ for all } k \in \{0, \dots, 3\},$$
$$\mathbf{w} \ge \mathbf{0},$$
(A.1)

where r_1, \ldots, r_n is a set of *n* grid locations spaced evenly over the interval $[0 \ \mu m, 500 \ \mu m]$, w_i is the number of particles with size equal to r_i , and $\mathbf{c} = (c_1, \ldots, c_n)$ is the cost vector whose components are defined by

$$c_{i} = \begin{cases} 1 & \text{if } r_{i} \in [85 \ \mu\text{m}, 150 \ \mu\text{m}], \\ 0 & \text{otherwise.} \end{cases} \quad \text{for all } i \in \{1, \dots, n\}.$$
(A.2)

The constraints ensure that the distribution has the specified moments, and the objective function counts the number of particles in the interval [85 μ m, 150 μ m].

Solving this LP with n = 10 grid locations gives $N_{\text{LP}}^U = 1.61 \times 10^3 \text{ cm}^{-3}$ as an "upper bound" on the number of particles in the interval. However, in Example 2.3.1, we have already shown that there exists a three-peaked distribution which satisfies the moment constraints and which has $1.85 \times 10^3 \text{ cm}^{-3}$ particles in the interval. This distribution exceeds the supposed "upper bound" by 13%. Thus, we see that N_{LP}^U is not an upper bound on the number of particles in the interval after all.

Now, one could reasonably object that we have used too few grid locations in the above argument. That's fair – we chose a small value of n to stress the point. As we

increase n and update the value of N_{LP}^U , the percentage by which the three-peaked distribution violates N_{LP}^U does decrease. However, even if we increase the number of grid locations, the *fundamental problem remains*. Unless the set of grid locations includes the points of support for the three-peaked distribution, there will always be a finite violation of the "bound".

This is true in general for the LP-based bounding method: unless the set of predetermined grid locations includes the (a priori unknown) points of support for the optimal distribution, the LP-based bound will be violated by this optimal distribution, and the violation will be finite. Intuitively, for a fine grid, the extent of the bound violation is likely to be small. But it is unclear how large a value of n is required before the violation is guaranteed to be negligibly small. So the "bound" isn't really a bound. It is at best an approximate bound, with unknown error. This is dissatisfying from a theoretical perspective. It is in this sense that the bounds provided by the LP method are not truly rigorous.

In contrast, the SDP-based bounding method is immune to this criticism, as it implicitly considers *all* distributions supported on \mathbb{R}_+ . There is no need to assume that the distribution is restricted to particles of particular sizes, r_1, \ldots, r_n .

A.0.3 Proof of Claim 1

Proof. Suppose we have a PSD described by a number density function f. By assumption, $f \in M^{\infty}(\mathbb{R}_+)$. Let $\tilde{\mu} = (\mu_0, \mu_1, \mu_2, ...)$ be the sequence of moments of the PSD, as defined by (2.2).

Pick an arbitrary $n \in \mathbb{N}$ and an arbitrary vector $\mathbf{p} = (p_0, p_1, ..., p_n) \in \mathbb{R}^{n+1}$. The vector \mathbf{p} can be used to define a polynomial:

$$p(x) = p_0 + p_1 x + \dots + p_n x^n = \sum_{j=0}^n p_j x^j.$$
 (A.3)

We then consider the integral $\int_0^{+\infty} p^2(x) f(x) dx$. Since $p^2(x) = (p(x))^2$ and f(x) are both nonnegative for each $x \in \mathbb{R}_+$, it follows that this integral is nonnegative. That is,

$$\int_0^{+\infty} p^2(x) f(x) dx \ge 0. \tag{A.4}$$

We can also expand the integral and relate it to the matrix $\mathbf{H}_n(\tilde{\mu})$ defined in Section 2.2.1:

$$\int_{0}^{+\infty} (p(x))^{2} f(x) dx = \int_{0}^{+\infty} \left(\sum_{j=0}^{n} p_{j} x^{j} \right)^{2} f(x) dx,$$

$$= \int_{0}^{+\infty} \sum_{i=0}^{n} \sum_{j=0}^{n} p_{i} p_{j} x^{i+j} f(x) dx,$$

$$= \sum_{i=0}^{n} \sum_{j=0}^{n} p_{i} p_{j} \int_{0}^{+\infty} x^{i+j} f(x) dx,$$

$$= \sum_{i=0}^{n} \sum_{j=0}^{n} p_{i} p_{j} \mu_{i+j},$$

$$= \mathbf{p}^{\mathrm{T}} \mathbf{H}_{n}(\tilde{\mu}) \mathbf{p}.$$

(A.5)

Together, (A.4) and (A.5) imply that $\mathbf{p}^{\mathrm{T}}\mathbf{H}_{n}(\tilde{\mu})\mathbf{p} \geq 0$. Since the choice of \mathbf{p} was arbitrary, we can conclude that $\mathbf{p}^{\mathrm{T}}\mathbf{H}_{n}(\tilde{\mu})\mathbf{p} \geq 0$ for any $\mathbf{p} \in \mathbb{R}^{n+1}$. Recalling the definition of a positive semidefinite matrix given in Section 2.2.1, we see that this is equivalent to the statement that $\mathbf{H}_{n}(\tilde{\mu}) \succeq \mathbf{0}$. Finally, since the choice of n was arbitrary, we see that we have $\mathbf{H}_{n}(\tilde{\mu}) \succeq \mathbf{0}$ for all $n \in \mathbb{N}$.

By a very similar argument, we can show that $\mathbf{B}_n(\tilde{\mu}) \succeq \mathbf{0}$ for all $n \in \mathbb{N}$. To reach this conclusion, one simply has to consider a slightly modified form of the above integral: $\int_0^{+\infty} x p^2(x) f(x) dx$ instead of $\int_0^{+\infty} p^2(x) f(x) dx$.

Technically, the above proof does not establish Claim 1 for all PSDs – only those PSDs which can be described by a number density function $f \in M^{\infty}(\mathbb{R}_+)$. In fact, there are some PSDs which cannot be described by a number density function. In particular, any PSD in which there are nonzero particles of any specific size (such as the three-peaked distribution appearing in Section 2.3.1) can only technically be described by a CDF or a measure. The general proof in terms of measures can be found in [42, Chapter 3]. This general proof requires more sophisticated notions of integration [67], which are beyond the scope of this paper. The above proof in terms of number density functions is sufficient to get the idea across and will be accessible to a larger audience, so we leave it at that.

Finally, we wish to point out that the above proof is readily extended to prove Claim 2. One simply has to consider the integral $\int_0^{+\infty} (x-a)(b-x)p^2(x)f(x)dx$.

A.0.4 Condition (2.7) implies nonnegativity of the moments

Proof. Pick an arbitrary $j \in \mathbb{N}$ such that j is odd. Pick any $n \in \mathbb{N}$ such that $2n+1 \ge j$. Condition (2.7) ensures that $\mathbf{B}_n(\tilde{\mu}) \succeq \mathbf{0}$ for this n. Using the characterization of positive semidefinite matrices given in (2.6), we can write

$$\mathbf{x}^{\mathrm{T}} \mathbf{B}_{n}(\tilde{\mu}) \mathbf{x} \ge 0, \text{ for all } \mathbf{x} \in \mathbb{R}^{n+1}.$$
 (A.6)

In particular, $\mathbf{e}_k^{\mathrm{T}} \mathbf{B}_n(\tilde{\mu}) \mathbf{e}_k \geq 0$, where k = (j+1)/2 and $\mathbf{e}_k \in \mathbb{R}^{n+1}$ is the unit vector, whose kth component is 1 and all other components are 0. Referring to the definition of $\mathbf{B}_n(\tilde{\mu})$ given in (2.4), we see that $\mathbf{e}_k^{\mathrm{T}} \mathbf{B}_n(\tilde{\mu}) \mathbf{e}_k = \mu_j$, and thus $\mu_j \geq 0$. In the case where j is even, the argument is much the same; we just use $\mathbf{H}_n(\tilde{\mu})$ in place of $\mathbf{B}_n(\tilde{\mu})$. Thus, we see that Condition (2.7) implies the nonnegativity of all moments.

A.0.5 Equivalence of Problems (2.10) and (2.11)

Proof. Let ρ_1^* denote the optimal value of Problem (2.10); similarly, let ρ_2^* denote the optimal value of Problem (2.11).

Consider an arbitrary feasible point of Problem (2.10), (g, h), with an associated objective value of ρ . Since Problems (2.11) and (2.10) have the same feasible set, (g, h) is also feasible for Problem (2.11). Decompose h into two functions h' and h''where

$$h'(x) = \begin{cases} 0, & \text{if } x \in [0, a), \\ h(x), & \text{if } x \in [a, b], \\ 0, & \text{if } x \in (b, +\infty), \end{cases}$$
(A.7)

and

$$h''(x) = \begin{cases} h(x), & \text{if } x \in [0, a), \\ 0, & \text{if } x \in [a, b], \\ h(x), & \text{if } x \in (b, +\infty). \end{cases}$$
(A.8)

By construction, h = h' + h''. Define a new point as $(\hat{g}, \hat{h}) \equiv (g + h', h'')$. Because of the linearity of integration, (\hat{g}, \hat{h}) is also feasible for Problem (2.11). Furthermore, its objective value is equal to

$$\begin{split} \int_{0}^{+\infty} \hat{g}(x) dx &= \int_{a}^{b} \hat{g}(x) dx, \\ &= \int_{a}^{b} \left(g(x) + h'(x) \right) dx, \\ &= \int_{a}^{b} \left(g(x) + h'(x) \right) dx + \int_{a}^{b} h''(x) dx, \\ &= \int_{a}^{b} g(x) dx + \int_{a}^{b} \left(h'(x) + h''(x) \right) dx, \\ &= \int_{a}^{b} g(x) dx + \int_{a}^{b} h(x) dx, \\ &= \int_{a}^{b} g(x) dx + \int_{a}^{b} h(x) dx, \\ &= \rho. \end{split}$$
(A.9)

In this way, for any feasible point of Problem (2.10), we can construct a feasible point for Problem (2.11) with the same objective value. This implies that $\rho_1^* \leq \rho_2^*$.

The argument to show that $\rho_1^* \ge \rho_2^*$ is very similar. In particular, given an arbitrary feasible point of Problem (2.11), (g, h), with an associated objective value of ρ , one can construct (\hat{g}, \hat{h}) which is feasible for Problem (2.10) and has an objective value of at least ρ .

A.0.6 Sufficiency of the *k*th order LMIs

Proof. By definition, the matrix $\mathbf{H}_k(\mu)$ is positive semidefinite if and only if $\mathbf{x}^T \mathbf{H}_k(\mu) \mathbf{x} \geq 0$ for all $\mathbf{x} \in \mathbb{R}^{k+1}$. As a special case, this implies that $\mathbf{x}^T \mathbf{H}_k(\mu) \mathbf{x} \geq 0$ for all vectors $\mathbf{x} \in \mathbb{R}^{k+1}$ such that $x_{k+1} = 0$. This is equivalent to the statement that $\mathbf{x}^T \mathbf{H}_{k-1}(\mu) \mathbf{x} \geq 0$ for all vectors $\mathbf{x} \in \mathbb{R}^k$. Thus, $\mathbf{H}_{k-1}(\mu)$ is positive semidefinite. By induction, $\mathbf{H}_n(\mu)$ is positive semidefinite for all n < k. Similar arguments can be

applied to the more complicated LMIs such as

$$(a+b)\mathbf{B}_k(\mu) - \mathbf{C}_k(\mu) - ab\mathbf{H}_k(\mu) \succeq \mathbf{0}.$$
 (A.10)

A.0.7 Derivation of lower bounding SDP

The problem for calculating the lower bound can be derived in much the same way as that for the upper bound. We start with an explicit statement of the quantity we would like to calculate:

$$N^{L} \equiv \min_{f} \quad \int_{a}^{b} f(x)dx$$

s.t. $f \in M^{\infty}(\mathbb{R}_{+}),$ (A.11)
 $\int_{0}^{+\infty} x^{j}f(x)dx = \mu_{j}, \quad \forall j \in \{0, ..., m\}.$

Then, we decompose f into three number density functions: $f = g + h_1 + h_2$, where h_2 is confined to [0, a], h_2 is confined to $[b, +\infty)$, and g, like f, is supported on \mathbb{R}_+ .

$$N^{L} = \min_{g,h_{1},h_{2}} \int_{a}^{b} g(x)dx + \int_{a}^{b} h_{1}(x)dx + \int_{a}^{b} h_{2}(x)dx$$

s.t. $g \in M^{\infty}(\mathbb{R}_{+}), h_{1} \in M^{\infty}([0,a]), h_{2} \in M^{\infty}([b,+\infty)),$
 $\int_{0}^{+\infty} x^{j}g(x)dx + \int_{0}^{+\infty} x^{j}h_{1}(x)dx + \int_{0}^{+\infty} x^{j}h_{2}(x)dx = \mu_{j},$
 $\forall j \in \{0,...,m\}.$ (A.12)

For the same reason that Problems (2.10) and (2.11) are equivalent, Problem (A.12) is equivalent to

$$N^{L} = \min_{g,h_{1},h_{2}} \int_{0}^{+\infty} g(x)dx$$

s.t. $g \in M^{\infty}(\mathbb{R}_{+}), h_{1} \in M^{\infty}([0,a]), h_{2} \in M^{\infty}([b,+\infty)),$
 $\int_{0}^{+\infty} x^{j}g(x)dx + \int_{0}^{+\infty} x^{j}h_{1}(x)dx + \int_{0}^{+\infty} x^{j}h_{2}(x)dx = \mu_{j},$
 $\forall j \in \{0,...,m\}.$ (A.13)

Next, we introduce the moments of g, h_1 , and h_2 :

$$N^{L} = \min_{g,h_{1},h_{2},\tilde{z},\tilde{y},\tilde{w}} z_{0}$$

s.t. $g \in M^{\infty}(\mathbb{R}_{+}), h_{1} \in M^{\infty}([0,a]), h_{2} \in M^{\infty}([b,+\infty)),$
 $z_{j} + y_{j} + w_{j} = \mu_{j}, \forall j \in \{0,...,m\},$
 $\int_{0}^{+\infty} x^{j}g(x)dx = z_{j}, \int_{0}^{+\infty} x^{j}h_{1}(x)dx = y_{j}, \int_{0}^{+\infty} x^{j}h_{2}(x)dx = w_{j},$
 $\forall j \in \mathbb{N}.$
(A.14)

Then, we add the necessary conditions for these moments:

$$N^{L} = \min_{g,h_{1},h_{2},\tilde{z},\tilde{y},\tilde{w}} z_{0}$$
s.t. $g \in M^{\infty}(\mathbb{R}_{+}), h_{1} \in M^{\infty}([0,a]), h_{2} \in M^{\infty}([b,+\infty)),$
 $z_{j} + y_{j} + w_{j} = \mu_{j}, \forall j \in \{0,...,m\},$
 $\int_{0}^{+\infty} x^{j}g(x)dx = z_{j}, \int_{0}^{+\infty} x^{j}h_{1}(x)dx = y_{j}, \int_{0}^{+\infty} x^{j}h_{2}(x)dx = w_{j},$
 $\forall j \in \mathbb{N},$
 $H_{n}(\tilde{z}) \succeq \mathbf{0}, \mathbf{B}_{n}(\tilde{z}) \succeq \mathbf{0}, \forall n \in \mathbb{N},$
 $H_{n}(\tilde{y}) \succeq \mathbf{0}, -\mathbf{C}_{n}(\tilde{y}) + a\mathbf{B}_{n}(\tilde{y}) \succeq \mathbf{0}, \forall n \in \mathbb{N},$
 $H_{n}(\tilde{w}) \succeq \mathbf{0}, \mathbf{B}_{n}(\tilde{w}) - b\mathbf{H}_{n}(\tilde{w}) \succeq \mathbf{0}, \forall n \in \mathbb{N}.$
(A.15)

The LMIs $\mathbf{B}_n(\tilde{w}) - b\mathbf{H}_n(\tilde{w}) \succeq \mathbf{0}$, for all $n \in \mathbb{N}$ can be derived in much the same way as the LMIs appearing in Claim 1. One simply considers the integral $\int_0^{+\infty} (x - w)^{-1} dw$

 $b)p^2(x)h_2(x)dx \ge 0.$

Next, we remove the functions, leaving just the moments:

$$\begin{split}
\tilde{N}^{L} &\equiv \min_{\tilde{z}, \tilde{y}, \tilde{w}} \quad z_{0} \\
\text{s.t.} \quad z_{j} + y_{j} + w_{j} = \mu_{j}, \quad \forall j \in \{0, ..., m\}, \\
\mathbf{H}_{n}(\tilde{z}) \succeq \mathbf{0}, \quad \mathbf{B}_{n}(\tilde{z}) \succeq \mathbf{0}, \quad \forall n \in \mathbb{N}, \\
\mathbf{H}_{n}(\tilde{y}) \succeq \mathbf{0}, \quad -\mathbf{C}_{n}(\tilde{y}) + a\mathbf{B}_{n}(\tilde{y}) \succeq \mathbf{0}, \quad \forall n \in \mathbb{N}, \\
\mathbf{H}_{n}(\tilde{w}) \succeq \mathbf{0}, \quad \mathbf{B}_{n}(\tilde{w}) - b\mathbf{H}_{n}(\tilde{w}) \succeq \mathbf{0}, \quad \forall n \in \mathbb{N}.
\end{split}$$
(A.16)

Finally, we truncate the LMIs to a finite k:

$$\underline{N}^{L} \equiv \min_{z,y,w} \qquad z_{0}$$
s.t. $z_{j} + y_{j} + w_{j} = \mu_{j}, \quad \forall j \in \{0, ..., m\},$

$$\mathbf{H}_{k}(z) \succeq \mathbf{0}, \quad \mathbf{B}_{k-1}(z) \succeq \mathbf{0}, \qquad (A.17)$$

$$\mathbf{H}_{k}(y) \succeq \mathbf{0}, \quad -\mathbf{C}_{k-1}(y) + a\mathbf{B}_{k-1}(y) \succeq \mathbf{0},$$

$$\mathbf{H}_{k}(w) \succeq \mathbf{0}, \quad \mathbf{B}_{k-1}(w) - b\mathbf{H}_{k-1}(w) \succeq \mathbf{0}.$$

A.0.8 Inconsistency of the "Dirac Delta Function"

The idea of a Dirac delta function, while it is a useful conceptual short-cut in some contexts, is inconsistent with the rest of mathematics.

To illustrate the problem, consider the usual definition of the Dirac delta function as the function which satisfies

$$\delta_z(x) = \begin{cases} 0, & \text{if } x \neq z, \\ +\infty, & \text{if } x = z, \end{cases}$$
(A.18)

and

$$\int_{-\infty}^{+\infty} \delta_z(x) dx = 1 \tag{A.19}$$

From Equation (A.18), we see that $\delta_z(x) = 0$ almost everywhere on the real line. In fact, the concept of a property holding *almost everywhere* is formally defined in mathematics. Moreover, it is well-established that the (standard) Lebesgue integral of a function that is zero almost everywhere is zero. This directly contradicts Equation (A.19). This is the first argument against the Dirac delta function.

There is also a second argument. Suppose we wanted to make use of the scaled Dirac delta function, $f = N\delta_z$ mentioned previously. How should it be defined? It is clear that we want

$$\int_{-\infty}^{+\infty} f(x)dx = \int_{-\infty}^{+\infty} N\delta_z(x)dx = N \int_{-\infty}^{+\infty} \delta_z(x)dx = N.$$
(A.20)

However, how does f assign values to the elements of \mathbb{R} ? We might say

$$f(x) = N\delta_z(x) = \begin{cases} N \cdot 0, & \text{if } x \neq z, \\ N \cdot +\infty, & \text{if } x = z, \end{cases} = \begin{cases} 0, & \text{if } x \neq z, \\ +\infty, & \text{if } x = z, \end{cases}$$
(A.21)

but then, comparing Equations (A.20) and (A.19), we see that f(x) is identical to $\delta_z(x)$ at every $x \in \mathbb{R}$. In other words, the functions are identical. This implies that integrating either should yield the same value, but this contradicts Equations (A.19) and (A.20).

While we cannot sensibly talk about Dirac delta *functions*, the concept of a Dirac *distribution* or Dirac *measure* is a perfectly well-defined, useful concept. See, for example [42].

A.0.9 Derivation of Problem (2.24)

While one can show directly how Problem (2.17) reduces to (2.24), the argument is cumbersome. It is actually easier to derive Problem (2.24) in much the same way we derived (2.17), starting from Problem (A.11). The only difference in the derivation is that there is no need to introduce the function h_1 supported on [0, a].
A.0.10 Proof that $x_2 < x_1$ implies $\overline{N}^U(x_2) \leq \overline{N}^U(x_1)$ and $\underline{N}^L(x_2) \leq \underline{N}^L(x_1)$

By assumption, both x_1 and x_2 are nonnegative.

We will first consider the lower bound inequality $\underline{N}^{L}(x_{2}) \leq \underline{N}^{L}(x_{1})$, because its proof is simpler.

Proof. Pick an arbitrary feasible point (z, w) of Problem (2.24) where $x = x_1$. By definition w must satisfy

$$\mathbf{H}_{k}(w) \succeq \mathbf{0} \text{ and } \mathbf{B}_{k-1}(w) - x_{1}\mathbf{H}_{k-1}(w) \succeq \mathbf{0}.$$
 (A.22)

As explained in Section A.0.6, the first LMI implies $\mathbf{H}_{k-1}(w) \succeq \mathbf{0}$. Since $x_2 < x_1$, this implies $-(x_2 - x_1)\mathbf{H}_{k-1}(w) \succeq \mathbf{0}$. Adding this to the second LMI, we obtain

$$\mathbf{B}_{k-1}(w) - x_1 \mathbf{H}_{k-1}(w) - (x_2 - x_1) \mathbf{H}_{k-1}(w) = \mathbf{B}_{k-1}(w) - x_2 \mathbf{H}_{k-1}(w) \succeq \mathbf{0}.$$
 (A.23)

Thus, (z, w) is also feasible for Problem (2.24) where $x = x_2$. Both problems have the same objective function, z_0 . Thus, we have $\underline{N}^L(x_2) \leq \underline{N}^L(x_1)$.

Next, we prove the upper bound inequality $\overline{N}^U(x_2) \leq \overline{N}^U(x_1)$.

Proof. Pick an arbitrary feasible point (y, z) of Problem (2.23) where $x = x_2$. By definition,

$$\mathbf{H}_{k}(z) \succeq \mathbf{0} \text{ and } x_{2}\mathbf{B}_{k-1}(z) - \mathbf{C}_{k-1}(z) \succeq \mathbf{0}.$$
 (A.24)

The matrix $\mathbf{C}_{k-1}(z)$ is nothing more than the $k \times k$ submatrix filling the bottom right corner of $\mathbf{H}_k(z)$ (i.e. the submatrix obtained by eliminating the first column and row of $\mathbf{H}_k(z)$). Thus, $\mathbf{H}_k(z) \succeq \mathbf{0}$ implies $\mathbf{C}_{k-1}(z) \succeq \mathbf{0}$.

Suppose that $\mathbf{B}_{k-1}(y)$ has a negative eigenvalue (i.e. that $\mathbf{B}_{k-1}(y) \not\geq \mathbf{0}$). This implies that there exists $\mathbf{v} \in \mathbb{R}^k$ such that $\mathbf{v}^{\mathrm{T}} \mathbf{B}_{k-1}(y) \mathbf{v} < 0$. Since $x_2 > 0$ and $-\mathbf{C}_{k-1}(z) \leq \mathbf{0}$, this implies

$$\mathbf{v}^{\mathrm{T}}\left(x_{2}\mathbf{B}_{k-1}(y) - \mathbf{C}_{k-1}(z)\right)\mathbf{v} < 0.$$
(A.25)

But this contradicts (A.24). Thus, our assumption that $\mathbf{B}_{k-1}(y) \succeq \mathbf{0}$ is false, and we have $\mathbf{B}_{k-1}(y) \succeq \mathbf{0}$.

Since $x_1 > x_2$, we have $(x_1 - x_2)\mathbf{B}_{k-1}(y) \succeq \mathbf{0}$. Adding this to the second LMI in (A.24), we obtain

$$(x_1 - x_2)\mathbf{B}_{k-1}(y) + x_2\mathbf{B}_{k-1}(z) - \mathbf{C}_{k-1}(z) = x_1\mathbf{B}_{k-1}(z) - \mathbf{C}_{k-1}(z) \succeq \mathbf{0}.$$
 (A.26)

Thus, (y, z) is also feasible for Problem (2.23) where $x = x_1$. Both problems have the same objective function, z_0 . Since (y, z) was an arbitrary feasible point of Problem (2.23) for $x = x_2$, this implies the desired inequality: $\overline{N}^U(x_2) \leq \overline{N}^U(x_1)$.

A.0.11 Algorithm for calculating a lower bound on the CDF

Algorithm 6 Generating a lower bound on the CDF Input:

- 1. Moments $\mu_0, ..., \mu_m$.
- 2. Estimate of the maximum particle size, x_{max} .
- 3. Number of test points, n.

Output: A function CDF_n^L which provides an lower bound for the true CDF on the interval $[0, x_{\text{max}}]$.

Algorithm:

Set $x_1 := 0$. Set $\ell_1 := x_1$. Solve Problem (2.24) to calculate $\underline{N}^L(x_1)$. Set $x_2 := x_{\max}$. Set $\ell_2 := x_2$. Solve Problem (2.24) to calculate $\underline{N}^L(x_2)$. Set $j^* := 1$. for i = 3, ..., n do Set $x_i := \frac{1}{2}(\ell_{j^*} + \ell_{j^*+1})$. Solve Problem (2.24) to calculate $\underline{N}^L(x_i)$. $\ell := \operatorname{sort}(\mathbf{x})$. Set $j^* \in \arg \max_{j \le i-1} \left\{ \left(\underline{N}^L(\ell_{j+1}) - \underline{N}^L(\ell_j) \right) (\ell_{j+1} - \ell_j) \right\}$ end for Set $\operatorname{CDF}_n^L(x) := \begin{cases} \underline{N}^L(\ell_j), & \text{if } x \in [\ell_j, \ell_{j+1}).\\ \underline{N}^L(\ell_n), & \text{if } x = x_{\max}. \end{cases}$ (A.27)

255

A.0.12 Algorithm for calculating a lower bound on $D\alpha$

Algorithm 7 Generating a lower bound on $D\alpha^L$

Input:

- 1. Moments $\mu_0, ..., \mu_m$.
- 2. The value of $\alpha \in [0, 100]$ for which you want to calculate bounds on $D\alpha$.
- 3. A tolerance $\epsilon > 0$.

Output: $\underline{D}\alpha_{\epsilon}^{L}$, which is an ϵ -close lower bound on $\overline{D}\alpha^{U}$. That is, it satisfies

$$\underline{\mathbf{D}}\alpha_{\epsilon}^{L} \leq \underline{\mathbf{D}}\alpha^{L} \leq \underline{\mathbf{D}}\alpha_{\epsilon}^{L} + \epsilon.$$
(A.28)

```
Algorithm:
Set \underline{D}\alpha_{\epsilon}^{L} := 0.
Set s := \mu_1 / \mu_0.
Solve Problem (2.34) to obtain \overline{\alpha}^U(s).
while \overline{\alpha}^U(s) < \alpha do
Set \underline{D\alpha}^L_{\epsilon} := s.
        Set s := 2s.
       Solve Problem (2.34) to obtain \overline{\alpha}^U(s).
end while
Set \overline{\mathrm{D}\alpha}^L_{\epsilon} := s
while \overline{D\alpha}_{\epsilon}^{L} - \underline{D\alpha}_{\epsilon}^{L} > \epsilon do
Set s := (\overline{D\alpha}_{\epsilon}^{L} + \underline{D\alpha}_{\epsilon}^{L})/2.
       Solve Problem (2.34) to obtain \overline{\alpha}^U(s).
       if \overline{\alpha}^U(s) \geq \alpha then
               Set \overline{\overline{\mathrm{D}\alpha}}^L := s.
        else
               Set D\alpha^L := s
        end if
end while
```

Appendix B

Supplemental Material for Chapter 3

The following supplemental material is intended to accompany Chapter 3. It is a slightly modified version of the supplemental material published along with the paper entitled "Bounds on Stochastic Chemical Kinetic Systems at Steady State" by Garrett R. Dowdy and Paul I. Barton.

B.0.1 A Moment Enumeration Scheme

To implement the methods described in this paper, it is necessary to be able to list systematically all moments up to a specified order. In particular, this is necessary for constructing the moment vectors $\boldsymbol{\mu}$, the **A** matrix, and the matrices $\mathbf{M}_n^{\mathbf{0}}(\boldsymbol{\mu})$, $\mathbf{M}_{n-1}^{\mathbf{0}}(\boldsymbol{\mu})$, and $\mathbf{M}_{n-1}^{\mathbf{e}_j}(\boldsymbol{\mu})$ (see Section 3.3.2). In the body of the paper, we took for granted that we could perform this enumeration, so as to not distract from the main message. However, enumerating the moments is not a trivial task. So, in this section, we describe how it can be done in some detail.

To enumerate the moments $\{\mu_{\mathbf{j}} : \mathbf{j} \in \mathbb{N}^{\hat{N}}\}\)$, we want to associate each multi-index $\mathbf{j} \in \mathbb{N}^{\hat{N}}\)$ with a single index $j \in \{1, 2, 3, ...\}\)$. To do this, we must pick a specific ordering for the multi-indices and moments.¹ In principle, there are many possible

¹We are well aware that we have already defined the word "order" for moments in Section 3.2.5, and that we are now introducing a close variant of the word with an entirely different meaning. This

orderings we could pick [10, Chapter 2]. However, so that we may take advantage of Smadbeck and Kaznessis's work on the efficient construction of the **A** matrices, we will use the same ordering that appears in that paper [63]. According to Smadbeck and Kaznessis, the moments should be grouped by order², and then "within each order the moments are simply indexed in descending order for each of the components". They offer the following example as the appropriate ordering of third-order moments of a system with three components:

$$\langle x_1^3 \rangle, \quad \langle x_1^2 x_2 \rangle, \quad \langle x_1^2 x_3 \rangle, \quad \langle x_1 x_2^2 \rangle, \quad \langle x_1 x_2 x_3 \rangle, \quad \langle x_1 x_3^2 \rangle, \quad \langle x_2^3 \rangle, \quad \langle x_2^2 x_3 \rangle, \quad \langle x_2 x_3^2 \rangle, \quad \langle x_3^3 \rangle.$$

$$(B.1)$$

Surprisingly, this ordering does not appear to have a well-established name. Taking inspiration from Smadbeck and Kaznessis's description, and the similar idea of "graded lexicographic order" [10, Chapter 2], we will call it *graded descending order*. While Smadbeck and Kaznessis's description and example may be sufficient to infer the pattern of the ordering, the idea is made more precise in the following definitions:

Definition 1 (Descending Order). Consider the two multi-indices $\mathbf{j} = (j_1, \ldots, j_{\hat{N}}), \mathbf{i} = (i_1, \ldots, i_{\hat{N}}) \in \mathbb{N}^{\hat{N}}$. We say that $\mathbf{j} <_d \mathbf{i}$ if the leftmost nonzero entry of the vector difference $\mathbf{j} - \mathbf{i} \in \mathbb{Z}^{\hat{N}}$ is positive. Furthermore, if $\mathbf{j} <_d \mathbf{i}$, we will write $\mu_{\mathbf{j}} <_d \mu_{\mathbf{i}}$.

This takes care of the "descending" idea in Smadbeck and Kaznessis's description. In particular, by this definition, we have

$$\langle x_1^3 \rangle <_d \langle x_1^2 x_2 \rangle <_d \langle x_1^2 x_3 \rangle <_d \langle x_1 x_2^2 \rangle <_d \langle x_1 x_2 x_3 \rangle <_d \langle x_1 x_3^2 \rangle <_d \langle x_2^3 \rangle <_d \cdots <_d \langle x_3^3 \rangle,$$

which is consistent with Smadbeck and Kaznessis's example.

To group the moments "by order", we build on Definition 1:

Definition 2 (Graded Descending Order). Consider the two multi-indices $\mathbf{j}, \mathbf{i} \in \mathbb{N}^{\hat{N}}$.

over-loading of the word is unfortunate, but it seems that "order" really is most appropriate word to express the present idea [10, Chapter 2], and the phrase "kth order moment" is well-established. As seen in the quote below, Smadbeck and Kaznessis encountered this difficulty too. We hope that the type of "order" we are talking about will be clear from the context.

² in the sense of Section 3.2.5

We say that $\mathbf{j} <_{gd} \mathbf{i}$ if

$$\mathbf{j}| < |\mathbf{i}|,\tag{B.2}$$

or if

$$|\mathbf{j}| = |\mathbf{i}|$$
 and $\mathbf{j} <_d \mathbf{i}$. (B.3)

Furthermore, if $\mathbf{j} <_{gd} \mathbf{i}$, we will write $\mu_{\mathbf{j}} <_{gd} \mu_{\mathbf{i}}$, and we will say that \mathbf{j} precedes \mathbf{i} .

The ordering relation " $\langle gd$ " allows us to enumerate the multi-indices $\{\mathbf{j} \in \mathbb{N}^{\hat{N}}\}$ in a well-defined sequence, associating each with a single index $j \in \{1, 2, 3, ...\}$. This association is shown in Table B.1. By construction, for each $j \in \{1, 2, 3, ...\}$, the multi-index in row j precedes the multi-index in row j + 1 in the sense of the order relation " $\langle gd$ ".

It is easy to verify that the multi-indices of Table B.1 are listed in graded descending order. However, we have the following practical question: given a multi-index $\mathbf{j} \in \mathbb{N}^{\hat{N}}$, how do we construct the next multi-index in the sequence, \mathbf{j}^* ? One solution is provided by Algorithm 8, below.

Algorithm 8 NextMultiIndex

Input: The current multi-index, $\mathbf{j} \in \mathbb{N}^{\hat{N}}$.

Output: The next multi-index in graded descending order, $\mathbf{j}^* \in \mathbb{N}^{\hat{N}}$.

Algorithm:

```
set j^* := j.

set i := \hat{N} - 1.

while i > 0 and j_i^* = 0 do

set i := i - 1.

end while

if i > 0 then

set j_i^* := j_i^* - 1.

end if

set j_{\hat{N}}^* := 0.

set j_{i+1}^* := 1 + j_{\hat{N}}.
```

It's straightforward to see that the \mathbf{j}^* produced by Algorithm 8 satisfies $\mathbf{j} <_{gd} \mathbf{j}^*$. There are only two cases to consider. First, if the condition of the "if" statement is not satisfied, then $|\mathbf{j}^*| = |\mathbf{j}| + 1$. By Definition 2, it follows that $\mathbf{j} <_{gd} \mathbf{j}^*$. Second, if

| Single Index | Multi-Index |
|---|-------------------------------|
| 1 | $(0,0,\ldots,0,0)$ |
| 2 | $(1,0,\ldots,0,0)$ |
| 3 | $(0,1,\ldots,0,0)$ |
| ÷ | ÷ |
| Ŵ | $(0,0,\ldots,1,0)$ |
| $\hat{N} + 1$ | $(0,0,\ldots,0,1)$ |
| $\hat{N}+2$ | $(2,0,\ldots,0,0)$ |
| $\hat{N}+3$ | $(1,1,\ldots,0,0)$ |
| ÷ | : |
| $2\hat{N}$ | $(1,0,\ldots,1,0)$ |
| $2\hat{N} + 1$ | $(1,0,\ldots,0,1)$ |
| $2\hat{N}+2$ | $(0,2,\ldots,0,0)$ |
| ÷ | : |
| $3\hat{N}-1$ | $(0,1,\ldots,1,0)$ |
| $3\hat{N}$ | $(0,1,\ldots,0,1)$ |
| ÷ | : |
| $\frac{1}{2}\hat{N}^2 + \frac{3}{2}\hat{N} - 1$ | $(0,0,\ldots,2,0)$ |
| $\frac{1}{2}\hat{N}^2 + \frac{3}{2}\hat{N}$ | $(0,\overline{0,\ldots,1},1)$ |
| $\frac{1}{2}\hat{N}^2 + \frac{3}{2}\hat{N} + 1$ | $(0,0,\ldots,0,2)$ |
| $\frac{1}{2}\hat{N}^2 + \frac{3}{2}\hat{N} + 2$ | $(3,\overline{0,\ldots,0,0})$ |
| : | : |

Table B.1: Association between single indices $(j \in \{1, 2, 3, ...\})$ and multi-indices $(\mathbf{j} \in \mathbb{N}^{\hat{N}})$

the "if" condition is satisfied, then $|\mathbf{j}^*| = |\mathbf{j}|$. Comparing the sub-indices, we see

$$j_{1}^{*} = j_{1},$$

 \vdots
 $j_{i-1}^{*} = j_{i-1},$
(B.4)

and $j_i^* = j_i - 1 < j_i$. So the leftmost nonzero entry of $\mathbf{j} - \mathbf{j}^*$ is positive, implying $\mathbf{j} <_d \mathbf{j}^*$. It follows that $\mathbf{j} <_{gd} \mathbf{j}^*$.

What's less obvious is that the \mathbf{j}^* produced by Algorithm 8 is necessarily the next multi-index in the sequence. In other words, we need to demonstrate that there can never be some $\mathbf{j}' \in \mathbb{N}^{\hat{N}}$ such that $\mathbf{j} <_{gd} \mathbf{j}' <_{gd} \mathbf{j}^*$.

Claim 5. Given a multi-index $\mathbf{j} \in \mathbb{N}^{\hat{N}}$, if we use Algorithm 8 to produce another multi-index $\mathbf{j}^* \in \mathbb{N}^{\hat{N}}$, there there exists no $\mathbf{j}' \in \mathbb{N}^{\hat{N}}$ such that $\mathbf{j} <_{gd} \mathbf{j}' <_{gd} \mathbf{j}^*$.

Proof. Assume that there exists some $\mathbf{j}' \in \mathbb{N}^{\hat{N}}$ such that $\mathbf{j} <_{gd} \mathbf{j}' <_{gd} \mathbf{j}^*$. We will show that this leads to contradiction. As before, there are two cases to consider:

- <u>Case 1</u>: $|\mathbf{j}^*| = |\mathbf{j}| + 1$. Examining Algorithm 8, we see that this case only ever occurs when $j_1 = \cdots = j_{\hat{N}-1} = 0$. It follows that $j_2^* = \cdots = j_{\hat{N}}^* = 0$. Furthermore, it follows that $|\mathbf{j}| = j_{\hat{N}}$ and $|\mathbf{j}^*| = j_1^*$. If $\mathbf{j} <_{gd} \mathbf{j}' <_{gd} \mathbf{j}^*$, then, by Definition 2, we must have either $|\mathbf{j}'| = |\mathbf{j}|$ or $|\mathbf{j}'| = |\mathbf{j}^*|$. We now consider these two sub-cases separately:
 - <u>Sub-case 1.1</u>: $|\mathbf{j}'| = |\mathbf{j}|$. If this equality holds, then, by Definition 2, we must have $\mathbf{j} <_d \mathbf{j}'$, meaning that the leftmost nonzero entry of $\mathbf{j} \mathbf{j}'$ is positive. Since $j_1 = \cdots = j_{\hat{N}-1} = 0$ and $\mathbf{j}' \in \mathbb{N}^{\hat{N}}$, this implies that $j'_1 = \cdots = j'_{\hat{N}-1} = 0$ and that $j'_{\hat{N}} < j_{\hat{N}}$. However, this implies that $|\mathbf{j}'| = j'_{\hat{N}} < j_{\hat{N}} = |\mathbf{j}|$, which contradicts our assumption that $|\mathbf{j}'| = |\mathbf{j}|$.
 - <u>Sub-case 1.2</u>: $|\mathbf{j}'| = |\mathbf{j}^*|$. If this equality holds, then, by Definition 2, we must have $\mathbf{j}' <_d \mathbf{j}^*$, meaning that the leftmost nonzero entry of $\mathbf{j}' \mathbf{j}^*$ is positive. This implies that $j'_1 \ge j^*_1$. If $j'_1 > j^*_1$, then $|\mathbf{j}'| \ge j'_1 > j^*_1 = |\mathbf{j}^*|$, which contradicts our assumption that $|\mathbf{j}'| = |\mathbf{j}^*|$. On the other hand, if

 $j'_1 = j^*_1$, then we must have $j'_i > 0$ for some $i \in \{2, \ldots, \hat{N}\}$, otherwise we would not have $\mathbf{j}' <_d \mathbf{j}^*$. It follows that $|\mathbf{j}'| \ge j'_1 + j'_i = j^*_1 + j'_i > j^*_1 = |\mathbf{j}^*|$, which contradicts our assumption that $|\mathbf{j}'| = |\mathbf{j}^*|$.

• <u>Case 2</u>: $|\mathbf{j}^*| = |\mathbf{j}|$. As noted previously, by the design of Algorithm 8, when this equality occurs, we have

$$j_{1}^{*} = j_{1},$$

 \vdots
 $j_{i-1}^{*} = j_{i-1},$
 $j_{i}^{*} = j_{i} - 1,$
(B.5)

for some $i \in \{1, \ldots, \hat{N} - 1\}$. Also, with $\mathbf{j} <_{gd} \mathbf{j}' <_{gd} \mathbf{j}^*$, the equality $|\mathbf{j}^*| = |\mathbf{j}|$ implies that $|\mathbf{j}'| = |\mathbf{j}^*| = |\mathbf{j}|$. By Definition 2, it follows that $\mathbf{j} <_d \mathbf{j}' <_d \mathbf{j}^*$. With (B.5), this implies that

$$j_{1}^{*} = j_{1}' = j_{1},$$

 \vdots
 $j_{i-1}^{*} = j_{i-1}' = j_{i-1},$
(B.6)

and that $j'_i = j_i$ or $j'_i = j^*_i$. We now consider these two sub-cases separately:

- <u>Sub-case 2.1</u>: $j'_i = j_i$. Inspecting Algorithm 8, we see that $j_{i+1} = \cdots = j_{\hat{N}-1} = 0$. With $\mathbf{j} <_d \mathbf{j}'$, this implies that $j'_{i+1} = \cdots = j'_{\hat{N}-1} = 0$ and that $j_{\hat{N}} > j'_{\hat{N}}$. However, this implies that

$$|\mathbf{j}| = \sum_{k=1}^{\hat{N}} j_k = j_{\hat{N}} + \sum_{k=1}^{\hat{N}-1} j_k = j_{\hat{N}} + \sum_{k=1}^{\hat{N}-1} j'_k > j'_{\hat{N}} + \sum_{k=1}^{\hat{N}-1} j'_k = |\mathbf{j}'|, \qquad (B.7)$$

which contradicts our prior conclusion that $|\mathbf{j}'| = |\mathbf{j}|$.

- <u>Sub-case 2.2</u>: $j'_i = j^*_i$. Inspecting Algorithm 8, we see that $j^*_{i+2} = \cdots = j^*_{\hat{N}} = 0$. With $\mathbf{j}' <_d \mathbf{j}^*$, this implies that $j'_{i+1} > j^*_{i+1}$ or that $j'_{i+1} = j^*_{i+1}$ and

that $j'_l > 0$ for some $l \in \{i + 2, ..., \hat{N}\}$. Either way, this implies

$$|\mathbf{j}'| = \sum_{k=1}^{\hat{N}} j'_k = \sum_{k=1}^{i} j'_k + \sum_{k=i+1}^{\hat{N}} j'_k = \sum_{k=1}^{i} j^*_k + \sum_{k=i+1}^{\hat{N}} j'_k > \sum_{k=1}^{i} j^*_k + \sum_{k=i+1}^{\hat{N}} j^*_k = |\mathbf{j}^*|,$$
(B.8)

which contradicts our prior conclusion that $|\mathbf{j}'| = |\mathbf{j}^*|$.

Thus, no matter which case or sub-case we consider, we always arrive at a contradiction. This completes the proof. $\hfill \Box$

Now that we have a well-defined and proven algorithm for constructing the next mult-index in the sequence, enumerating the multi-indices $\{\mathbf{j} \in \mathbb{N}^{\hat{N}}\}$ is trivial. Furthermore, we can define a simple function which maps single indices to multi-indices, given by Algorithm 9.

Algorithm 9 MapSingleToMultiIndex Input:

- 1. A single index, $j \in \{1, 2, 3, ...\}$.
- 2. The number of independent components, $\hat{N} \in \{1, 2, 3, ...\}$.

Output: The multi-index $\mathbf{j} \in \mathbb{N}^{\hat{N}}$ corresponding to the single index j.

```
Algorithm:
set \mathbf{j} := (0, ..., 0) \in \mathbb{N}^{\hat{N}}.
for i = 2 to j do
set \mathbf{j} := \text{NextMultiIndex}(\mathbf{j}).
end for
```

In what follows "MapSingleToMultiIndex (j, \hat{N}) " will be abbreviated as " $I_{s \to m}(j, \hat{N})$ ", where the "I" stands for "index", "s" stands for "single" and "m" stands for "multi".

B.0.2 Definition of Matrices $\mathbf{M}_{n}^{\mathbf{0}}(\boldsymbol{\mu})$, $\mathbf{M}_{n-1}^{\mathbf{0}}(\boldsymbol{\mu})$, and $\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\boldsymbol{\mu})$

The notation $\mathbf{M}_{k}^{\mathbf{v}}(\boldsymbol{\mu})$ has three parts:

- $k \in \mathbb{N}$, which specifies the size of the matrix,
- $\mathbf{v} \in \mathbb{N}^{\hat{N}}$, which specifies a multi-index "offset";

• and μ , the vector from which the entries of the matrix will be drawn.

In particular, the matrix $\mathbf{M}_{k}^{\mathbf{v}}(\boldsymbol{\mu})$ will have $\binom{k+\hat{N}}{k}$ rows and columns, and the entry in the *i*th row and *j*th column is given by

$$(\mathbf{M}_{k}^{\mathbf{v}}(\boldsymbol{\mu}))_{i,j} \equiv \boldsymbol{\mu}_{\mathbf{v}+I_{s\to m}(i,\hat{N})+I_{s\to m}(j,\hat{N})}, \quad \forall (i,j) \in \left\{1,\dots, \binom{k+\hat{N}}{k}\right\}^{2}.$$
 (B.9)

The function $I_{s\to m}(\cdot, \hat{N})$ converts the row and column (single) indices into multiindices. These two multi-indices are then added together and offset by **v**.

This definition of $\mathbf{M}_{k}^{\mathbf{v}}(\boldsymbol{\mu})$ has the virtue of being very precise. However, it is somewhat difficult to interpret. An equivalent way to define $\mathbf{M}_{k}^{\mathbf{v}}(\boldsymbol{\mu})$ is to say that the rows and columns are indexed by all multi-indices $\mathbf{i}, \mathbf{j} \in \mathbb{N}^{\hat{N}}$ such that $|\mathbf{i}|, |\mathbf{j}| \leq k$, where the multi-indices appear in graded descending order. Then, the element of the matrix in the row corresponding to \mathbf{i} and the column corresponding to \mathbf{j} is given by

$$(\mathbf{M}_{k}^{\mathbf{v}}(\boldsymbol{\mu}))_{\mathbf{i},\mathbf{j}} \equiv \mu_{\mathbf{v}+\mathbf{i}+\mathbf{j}}, \quad \forall \mathbf{i}, \mathbf{j} \in \mathbb{N}^{\hat{N}} \text{ s.t. } |\mathbf{i}|, |\mathbf{j}| \le k.$$
 (B.10)

The number of rows and columns is equal to $\binom{k+\hat{N}}{k}$ because this is precisely the number of multi-indices (and moments) with order at most k. This can be shown through the classic "bars and stars" argument from combinatorics.

The above definition of the generic matrix $\mathbf{M}_{k}^{\mathbf{v}}(\boldsymbol{\mu})$ can be used to construct the specific matrices $\mathbf{M}_{n}^{\mathbf{0}}(\boldsymbol{\mu})$, $\mathbf{M}_{n-1}^{\mathbf{0}}(\boldsymbol{\mu})$, and $\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\boldsymbol{\mu})$. The structure of these matrices is closely tied to the derivation of the LMIs (3.27) - (3.29), which will be explored in the next section.

B.0.3 Derivation of LMIs

In this section, we derive the LMIs (3.27) - (3.29). In doing so, we assume that the reader has read and understood Sections B.0.2 and B.0.1.

We will begin with the most basic of the three: $\mathbf{M}_n^{\mathbf{0}}(\boldsymbol{\mu}) \succeq \mathbf{0}$. To prove this

inequality, consider a polynomial

$$p(\hat{\mathbf{x}}) \equiv \sum_{\mathbf{j} \in \mathbb{N}^{\hat{N}}}^{|\mathbf{j}| \le n} p_{\mathbf{j}} \hat{\mathbf{x}}^{\mathbf{j}}$$
(B.11)

where the sum is over all multi-indices $\mathbf{j} \in \mathbb{N}^{\hat{N}}$ with order at most $n \geq 1$, and the coefficients $\{p_{\mathbf{j}}\}_{\mathbf{j}\in\mathbb{N}^{\hat{N}}}^{|\mathbf{j}|\leq n} \subset \mathbb{R}$ are chosen arbitrarily. If we take the square of this polynomial, $p(\hat{\mathbf{x}})^2$, the result is a function which is nonnegative for all $\hat{\mathbf{x}} \in \mathbb{N}^{\hat{N}}$, and, in particular, for all $\hat{\mathbf{x}} \in \hat{\mathcal{X}}$ (since $\hat{\mathcal{X}}$ is a subset of $\mathbb{N}^{\hat{N}}$). Furthermore, since any steady-state distribution $P_{\rm ss}(\cdot)$ is a *probability* distribution over $\hat{\mathcal{X}}$, we also have $P_{\rm ss}(\hat{\mathbf{x}}) \geq 0$, for all $\hat{\mathbf{x}} \in \hat{\mathcal{X}}$. It follows that

$$\sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}} p(\hat{\mathbf{x}})^2 P_{\rm ss}(\hat{\mathbf{x}}) \ge 0.$$
(B.12)

Expanding the squared polynomial, we obtain

$$\begin{split} \sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}} p(\hat{\mathbf{x}})^2 P_{\mathrm{ss}}(\hat{\mathbf{x}}) &= \sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}} \left(\sum_{\mathbf{j}\in\mathbb{N}^{\hat{N}}}^{|\mathbf{j}|\leq n} p_{\mathbf{j}} \hat{\mathbf{x}}^{\mathbf{j}} \right)^2 P_{\mathrm{ss}}(\hat{\mathbf{x}}), \\ &= \sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}} \left(\sum_{\mathbf{j}\in\mathbb{N}^{\hat{N}}}^{|\mathbf{j}|\leq n} \sum_{\mathbf{i}\in\mathbb{N}^{\hat{N}}}^{|\mathbf{i}|\leq n} p_{\mathbf{j}} p_{\mathbf{i}} \hat{\mathbf{x}}^{\mathbf{j}} \hat{\mathbf{x}}^{\mathbf{j}} \right) P_{\mathrm{ss}}(\hat{\mathbf{x}}), \\ &= \sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}} \left(\sum_{\mathbf{j}\in\mathbb{N}^{\hat{N}}}^{|\mathbf{j}|\leq n} \sum_{\mathbf{i}\in\mathbb{N}^{\hat{N}}}^{|\mathbf{i}|\leq n} p_{\mathbf{j}} p_{\mathbf{i}} \hat{\mathbf{x}}^{\mathbf{j}+\mathbf{i}} \right) P_{\mathrm{ss}}(\hat{\mathbf{x}}), \\ &= \sum_{\mathbf{j}\in\mathbb{N}^{\hat{N}}}^{|\mathbf{j}|\leq n} \sum_{\mathbf{i}\in\mathbb{N}^{\hat{N}}}^{|\mathbf{i}|\leq n} p_{\mathbf{j}} p_{\mathbf{i}} \left(\sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}} \hat{\mathbf{x}}^{\mathbf{j}+\mathbf{i}} P_{\mathrm{ss}}(\hat{\mathbf{x}}) \right), \\ &= \sum_{\mathbf{j}\in\mathbb{N}^{\hat{N}}}^{|\mathbf{j}|\leq n} \sum_{\mathbf{i}\in\mathbb{N}^{\hat{N}}}^{|\mathbf{i}|\leq n} p_{\mathbf{j}} p_{\mathbf{i}} \mu_{\mathbf{j}+\mathbf{i}}. \end{split}$$
(B.13)

Combining this result with Inequality (B.12), we have

$$\sum_{\mathbf{j}\in\mathbb{N}^{\hat{N}}}^{|\mathbf{j}|\leq n} \sum_{\mathbf{i}\in\mathbb{N}^{\hat{N}}}^{|\mathbf{i}|\leq n} p_{\mathbf{j}} p_{\mathbf{i}} \mu_{\mathbf{j}+\mathbf{i}} \ge 0.$$
(B.14)

Now, let $\mathbf{p} \in \mathbb{R}^{\binom{\hat{N}+n}{n}}$ be the vector obtained by listing the coefficients $\{p_{\mathbf{j}}\}_{\mathbf{j}\in\mathbb{N}^{\hat{N}}}^{|\mathbf{j}|\leq n} \subset \mathbb{R}$ so that their subscripts are in graded descending order. Then, by design, the left-hand side of the above inequality is equal to $\mathbf{p}^{\mathrm{T}}\mathbf{M}_{n}^{\mathbf{0}}(\boldsymbol{\mu})\mathbf{p}$. (We don't expect this to be obvious, but it is critical; convince yourself before proceeding.) What we're left with, then, is

$$\mathbf{p}^{\mathrm{T}}\mathbf{M}_{n}^{\mathbf{0}}(\boldsymbol{\mu})\mathbf{p} \geq 0.$$
(B.15)

Recall that our choice of the coefficients $\{p_{\mathbf{j}}\}_{\mathbf{j}\in\mathbb{N}^{\hat{N}}}^{|\mathbf{j}|\leq n} \subset \mathbb{R}$ was arbitrary. This means that Inequality (B.15) holds for all vectors $\mathbf{p}\in\mathbb{R}^{\binom{\hat{N}+n}{n}}$. As noted in Section III B, this is equivalent to the statement that $\mathbf{M}_{n}^{\mathbf{0}}(\boldsymbol{\mu})\succeq \mathbf{0}^{.3}$ This completes the derivation of LMI (3.27).

The derivation of LMI (3.28) is similar, but has one novel aspect. Again, we consider an arbitrary polynomial

$$p(\hat{\mathbf{x}}) \equiv \sum_{\mathbf{j} \in \mathbb{N}^{\hat{N}}}^{|\mathbf{j}| \le n-1} p_{\mathbf{j}} \hat{\mathbf{x}}^{\mathbf{j}}.$$
 (B.16)

Then, we pick an arbitrary independent species $j \in \{1, \ldots, \hat{N}\}$ and consider the sum

$$\sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}} \hat{x}_j p(\hat{\mathbf{x}})^2 P_{\rm ss}(\hat{\mathbf{x}}). \tag{B.17}$$

As before, $p(\hat{\mathbf{x}})^2$ and $P_{ss}(\hat{\mathbf{x}})$ are nonnegative over $\hat{\mathcal{X}}$. Furthermore, since \hat{x}_j is the molecular count of species j, we have $\hat{x}_j \geq 0$. It follows that

$$\sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}} \hat{x}_j p(\hat{\mathbf{x}})^2 P_{\rm ss}(\hat{\mathbf{x}}) \ge 0.$$
(B.18)

³Notice also that our choice of $n \in \mathbb{N}$ was arbitrary, so this LMI actually holds for any $n \in \mathbb{N}$.

As before, we expand the squared polynomial and obtain

$$\sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}} \hat{x}_{j} p(\hat{\mathbf{x}})^{2} P_{\mathrm{ss}}(\hat{\mathbf{x}}) = \sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}} \hat{x}_{j} \left(\sum_{\mathbf{j}\in\mathbb{N}^{\hat{N}}}^{|\mathbf{j}|\leq n-1} p_{\mathbf{j}} \hat{\mathbf{x}}^{\mathbf{j}} \right)^{2} P_{\mathrm{ss}}(\hat{\mathbf{x}}),$$

$$= \sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}} \hat{x}_{j} \left(\sum_{\mathbf{j}\in\mathbb{N}^{\hat{N}}}^{|\mathbf{j}|\leq n-1} p_{\mathbf{j}} p_{\mathbf{j}} \hat{\mathbf{x}}^{\mathbf{j}} \hat{\mathbf{x}}^{\mathbf{j}} \right) P_{\mathrm{ss}}(\hat{\mathbf{x}}),$$

$$= \sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}} \left(\sum_{\mathbf{j}\in\mathbb{N}^{\hat{N}}}^{|\mathbf{j}|\leq n-1} p_{\mathbf{j}} p_{\mathbf{j}} \hat{\mathbf{x}}^{\mathbf{j}+\mathbf{i}+\mathbf{e}_{j}} \right) P_{\mathrm{ss}}(\hat{\mathbf{x}}),$$

$$= \sum_{\mathbf{j}\in\mathbb{N}^{\hat{N}}}^{|\mathbf{j}|\leq n-1} \sum_{\mathbf{i}\in\mathbb{N}^{\hat{N}}}^{|\mathbf{i}|\leq n-1} p_{\mathbf{j}} p_{\mathbf{i}} \left(\sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}} \hat{\mathbf{x}}^{\mathbf{j}+\mathbf{i}+\mathbf{e}_{j}} P_{\mathrm{ss}}(\hat{\mathbf{x}}) \right),$$

$$= \sum_{\mathbf{j}\in\mathbb{N}^{\hat{N}}}^{|\mathbf{j}|\leq n-1} \sum_{\mathbf{i}\in\mathbb{N}^{\hat{N}}}^{|\mathbf{i}|\leq n-1} p_{\mathbf{j}} p_{\mathbf{i}} \mu_{\mathbf{j}+\mathbf{i}+\mathbf{e}_{j}}.$$
(B.19)

With Inequality (B.18), it follows that

$$\sum_{\mathbf{j}\in\mathbb{N}^{\hat{N}}}^{|\mathbf{j}|\leq n-1}\sum_{\mathbf{i}\in\mathbb{N}^{\hat{N}}}^{|\mathbf{i}|\leq n-1}p_{\mathbf{j}}p_{\mathbf{i}}\mu_{\mathbf{j}+\mathbf{i}+\mathbf{e}_{j}}\geq 0.$$
(B.20)

Now, let $\mathbf{p} \in \mathbb{R}^{\binom{\hat{N}+n-1}{n-1}}$ be the vector obtained by listing the coefficients $\{p_{\mathbf{j}}\}_{\mathbf{j}\in\mathbb{N}^{\hat{N}}}^{|\mathbf{j}|\leq n-1} \subset \mathbb{R}$ so that their subscripts are in graded descending order. Then, by design, the lefthand side of the above inequality is equal to $\mathbf{p}^{\mathrm{T}}\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\boldsymbol{\mu})\mathbf{p}$. Since the vector \mathbf{p} and the index $j \in \{1, \ldots, \hat{N}\}$ were arbitrary, it follows that

$$\mathbf{M}_{n-1}^{\mathbf{e}_j}(\boldsymbol{\mu}) \succeq \mathbf{0}, \quad \forall j \in \{1, \dots, \hat{N}\}.$$
(B.21)

The derivation of LMI (3.29) is so similar we will not repeat the details. The key difference from the preceding derivation is that the inequality for the molecular count of dependent species $k \in \{1, \ldots, L\}$,

$$\tilde{x}_{0,k} + \sum_{j=1}^{\hat{N}} \beta_{k,j} (\hat{x}_{0,j} - \hat{x}_j) \ge 0, \qquad (B.22)$$

is used in place of the inequality $\hat{x}_j \ge 0$. Given the preceding derivations, the con-

nection between (B.22) and

$$\left(\tilde{x}_{0,k} + \sum_{j=1}^{\hat{N}} \beta_{k,j} \hat{x}_{0,j}\right) \mathbf{M}_{n-1}^{\mathbf{0}}(\boldsymbol{\mu}) - \sum_{j=1}^{\hat{N}} \beta_{k,j} \mathbf{M}_{n-1}^{\mathbf{e}_j}(\boldsymbol{\mu}) \succeq \mathbf{0}, \quad \forall k \in \{1, \dots, L\}.$$
(B.23)

should be reasonably clear.

Note that, in each of the above derivations, the value of the parameter $n \in \mathbb{N}$ was unimportant, so long as it satisfies the inequality $n \geq 1$. This means that the LMIs we have derived are valid for any $n \in \mathbb{N}$ such that $n \geq 1$. This naturally leads to the question of what value of n to choose in construction of our SDPs. In general, the greater the value of n, the more we will restrict the feasible set of the SDPs. This is desirable, since it may lead to tighter bounds. However, there is a trade-off, as a larger value of n implies larger matrices. We will explain our method of picking n in Section B.0.5.

B.0.4 Proof that LMIs (3.27) - (3.29) Imply (3.24) - (3.26)

It is simplest to prove that (3.28) implies (3.24), so we will start there.

Claim 6. If

$$\mathbf{M}_{n-1}^{\mathbf{e}_j}(\boldsymbol{\mu}) \succeq \mathbf{0}, \quad \forall j \in \{1, \dots, \hat{N}\},\$$

for any $n \ge 1$, then

$$\mu_{\mathbf{e}_i} \ge 0, \quad \forall j \in \{1, \dots, \hat{N}\}. \tag{B.24}$$

Proof. Pick any $j \in \{1, \ldots, \hat{N}\}$. Recall that $\mathbf{M}_{n-1}^{\mathbf{e}_j}(\boldsymbol{\mu}) \succeq \mathbf{0}$ is equivalent to the statement that

$$\mathbf{p}\mathbf{M}_{n}^{\mathbf{e}_{j}}(\boldsymbol{\mu})\mathbf{p} \geq 0, \quad \forall \mathbf{p} \in \mathbb{R}^{\binom{\hat{N}+n-1}{n-1}}.$$

In particular,

$$\mathbf{e}_1 \mathbf{M}_n^{\mathbf{e}_j}(\boldsymbol{\mu}) \mathbf{e}_1 \ge 0.$$

The left-hand side of this inequality is equal to the entry in the first row and first

column of the matrix $\mathbf{M}_n^{\mathbf{e}_j}(\boldsymbol{\mu})$. From Definition (B.9) , we see that this is equal to

$$(\mathbf{M}_{n}^{\mathbf{e}_{j}}(\boldsymbol{\mu}))_{1,1} = \mu_{\mathbf{e}_{j}+I_{s\to m}(1,\hat{N})+I_{s\to m}(1,\hat{N})} = \mu_{\mathbf{e}_{j}+\mathbf{0}+\mathbf{0}} = \mu_{\mathbf{e}_{j}}$$
(B.25)

It follows that $\mu_{\mathbf{e}_j} \geq 0$. Since our choice of j was arbitrary, this is true for all $j \in \{1, \ldots, \hat{N}\}$.

The proof that (3.29) implies (3.25) is similar.

Claim 7. If

$$\left(\tilde{x}_{0,k} + \sum_{j=1}^{\hat{N}} \beta_{k,j} \hat{x}_{0,j}\right) \mathbf{M}_{n-1}^{\mathbf{0}}(\boldsymbol{\mu}) - \sum_{j=1}^{\hat{N}} \beta_{k,j} \mathbf{M}_{n-1}^{\mathbf{e}_j}(\boldsymbol{\mu}) \succeq \mathbf{0}, \quad \forall k \in \{1, \dots, L\},$$

for any $n \geq 1$ and $\mu_0 = 1$, then

$$\left(\tilde{x}_{0,k} + \sum_{j=1}^{\hat{N}} \beta_{k,j} \hat{x}_{0,j}\right) - \sum_{j=1}^{\hat{N}} \beta_{k,j} \mu_{\mathbf{e}_j} \ge 0, \quad \forall k \in \{1, \dots, L\}$$
(B.26)

Proof. Pick any $k \in \{1, \ldots, L\}$. Recall that

$$\left(\tilde{x}_{0,k} + \sum_{j=1}^{\hat{N}} \beta_{k,j} \hat{x}_{0,j}\right) \mathbf{M}_{n-1}^{\mathbf{0}}(\boldsymbol{\mu}) - \sum_{j=1}^{\hat{N}} \beta_{k,j} \mathbf{M}_{n-1}^{\mathbf{e}_j}(\boldsymbol{\mu}) \succeq \mathbf{0}$$
(B.27)

is equivalent to the statement that

$$\mathbf{p}^{\mathrm{T}}\left(\left(\tilde{x}_{0,k} + \sum_{j=1}^{\hat{N}} \beta_{k,j} \hat{x}_{0,j}\right) \mathbf{M}_{n-1}^{\mathbf{0}}(\boldsymbol{\mu}) - \sum_{j=1}^{\hat{N}} \beta_{k,j} \mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\boldsymbol{\mu})\right) \mathbf{p} \ge 0$$
(B.28)

for all $\mathbf{p} \in \mathbb{R}^{\binom{\hat{N}+n-1}{n-1}}$. In particular,

$$0 \leq \mathbf{e}_{1}^{\mathrm{T}} \left(\left(\tilde{x}_{0,k} + \sum_{j=1}^{\hat{N}} \beta_{k,j} \hat{x}_{0,j} \right) \mathbf{M}_{n-1}^{\mathbf{0}}(\boldsymbol{\mu}) - \sum_{j=1}^{\hat{N}} \beta_{k,j} \mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\boldsymbol{\mu}) \right) \mathbf{e}_{1}$$

$$= \left(\tilde{x}_{0,k} + \sum_{j=1}^{\hat{N}} \beta_{k,j} \hat{x}_{0,j} \right) \mathbf{e}_{1}^{\mathrm{T}} \mathbf{M}_{n-1}^{\mathbf{0}}(\boldsymbol{\mu}) \mathbf{e}_{1} - \sum_{j=1}^{\hat{N}} \beta_{k,j} \mathbf{e}_{1}^{\mathrm{T}} \mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\boldsymbol{\mu}) \mathbf{e}_{1}$$
(B.29)

We've already seen that $\mathbf{e}_1^{\mathrm{T}} \mathbf{M}_{n-1}^{\mathbf{e}_j}(\boldsymbol{\mu}) \mathbf{e}_1 = \mu_{\mathbf{e}_j}$. Similarly, $\mathbf{e}_1^{\mathrm{T}} \mathbf{M}_{n-1}^{\mathbf{0}}(\boldsymbol{\mu}) \mathbf{e}_1 = \mu_{\mathbf{0}}$. Substi-

tuting these moments into the above inequality, we obtain

$$\left(\tilde{x}_{0,k} + \sum_{j=1}^{\hat{N}} \beta_{k,j} \hat{x}_{0,j}\right) \mu_{\mathbf{0}} - \sum_{j=1}^{\hat{N}} \beta_{k,j} \mu_{\mathbf{e}_j} \ge 0.$$

Finally, using $\mu_0 = 1$ we obtain Inequality (B.26).

The proof that LMI (3.27) implies (3.26) has a different flavor.

Claim 8. If $\mathbf{M}^{\mathbf{0}}_n(\boldsymbol{\mu}) \succeq \mathbf{0}$ for any $n \geq 1$ and $\mu_{\mathbf{0}} = 1$, then

$$\begin{bmatrix} \mu_{2\mathbf{e}_{1}} - \mu_{\mathbf{e}_{1}}^{2} & \mu_{\mathbf{e}_{1}+\mathbf{e}_{2}} - \mu_{\mathbf{e}_{1}}\mu_{\mathbf{e}_{2}} & \dots & \mu_{\mathbf{e}_{1}+\mathbf{e}_{\hat{N}}} - \mu_{\mathbf{e}_{1}}\mu_{\mathbf{e}_{\hat{N}}} \\ \mu_{\mathbf{e}_{2}+\mathbf{e}_{1}} - \mu_{\mathbf{e}_{2}}\mu_{\mathbf{e}_{1}} & \mu_{2\mathbf{e}_{2}} - \mu_{\mathbf{e}_{2}}^{2} & \dots & \mu_{\mathbf{e}_{2}+\mathbf{e}_{\hat{N}}} - \mu_{\mathbf{e}_{2}}\mu_{\mathbf{e}_{\hat{N}}} \\ \vdots & \vdots & \ddots & \vdots \\ \mu_{\mathbf{e}_{\hat{N}}+\mathbf{e}_{1}} - \mu_{\mathbf{e}_{\hat{N}}}\mu_{\mathbf{e}_{1}} & \mu_{\mathbf{e}_{\hat{N}}+\mathbf{e}_{2}} - \mu_{\mathbf{e}_{\hat{N}}}\mu_{\mathbf{e}_{2}} & \dots & \mu_{2\mathbf{e}_{\hat{N}}} - \mu_{\mathbf{e}_{\hat{N}}}^{2} \end{bmatrix} \succeq \mathbf{0}.$$
(B.30)

Proof. If $\mathbf{M}_n^{\mathbf{0}}(\boldsymbol{\mu}) \succeq \mathbf{0}$ for any $n \ge 1$, it follows that $\mathbf{M}_1^{\mathbf{0}}(\boldsymbol{\mu}) \succeq \mathbf{0}$. This is true because $\mathbf{M}_1^{\mathbf{0}}(\boldsymbol{\mu})$ is a leading principal sub-matrix of $\mathbf{M}_n^{\mathbf{0}}(\boldsymbol{\mu})$, and any principal sub-matrix of a positive semidefinite matrix is necessarily positive semidefinite. Using Definition (B.9) or (B.10), we can write the LMI $\mathbf{M}_1^{\mathbf{0}}(\boldsymbol{\mu}) \succeq \mathbf{0}$ explicitly as

$$\mathbf{M}_{1}^{\mathbf{0}}(\boldsymbol{\mu}) = \begin{bmatrix} \mu_{\mathbf{0}} & \mu_{\mathbf{e}_{1}} & \dots & \mu_{\mathbf{e}_{\hat{N}}} \\ \mu_{\mathbf{e}_{1}} & \mu_{2\mathbf{e}_{1}} & \dots & \mu_{\mathbf{e}_{1}+\mathbf{e}_{\hat{N}}} \\ \vdots & \vdots & \ddots & \vdots \\ \mu_{\mathbf{e}_{\hat{N}}} & \mu_{\mathbf{e}_{\hat{N}}+\mathbf{e}_{1}} & \dots & \mu_{2\mathbf{e}_{\hat{N}}} \end{bmatrix} \succeq \mathbf{0}.$$
(B.31)

Applying the Schur Complement Lemma, this is equivalent to

$$\begin{bmatrix} \mu_{2\mathbf{e}_{1}} & \dots & \mu_{\mathbf{e}_{1}+\mathbf{e}_{\hat{N}}} \\ \vdots & \ddots & \vdots \\ \mu_{\mathbf{e}_{\hat{N}}+\mathbf{e}_{1}} & \dots & \mu_{2\mathbf{e}_{\hat{N}}} \end{bmatrix} - \begin{bmatrix} \mu_{\mathbf{e}_{1}} \\ \vdots \\ \mu_{\mathbf{e}_{\hat{N}}} \end{bmatrix} \frac{1}{\mu_{\mathbf{0}}} \begin{bmatrix} \mu_{\mathbf{e}_{1}} & \dots & \mu_{\mathbf{e}_{\hat{N}}} \end{bmatrix} \succeq \mathbf{0}.$$
(B.32)

Using the fact that $\mu_0 = 1$ and expanding the second product, we obtain

$$\begin{bmatrix} \mu_{2\mathbf{e}_{1}} & \dots & \mu_{\mathbf{e}_{1}+\mathbf{e}_{\hat{N}}} \\ \vdots & \ddots & \vdots \\ \mu_{\mathbf{e}_{\hat{N}}+\mathbf{e}_{1}} & \dots & \mu_{2\mathbf{e}_{\hat{N}}} \end{bmatrix} - \begin{bmatrix} \mu_{\mathbf{e}_{1}}^{2} & \dots & \mu_{\mathbf{e}_{1}}\mu_{\mathbf{e}_{\hat{N}}} \\ \vdots & \ddots & \vdots \\ \mu_{\mathbf{e}_{\hat{N}}}\mu_{\mathbf{e}_{1}} & \dots & \mu_{\mathbf{e}_{\hat{N}}}^{2} \end{bmatrix} \succeq \mathbf{0}.$$
(B.33)

Combining the two matrices into one, we obtain LMI (B.30). $\hfill \square$

B.0.5 Size of the SDPs

At this point, we have most of information necessary to compute the measures of SDP size appearing in Table 3.1. The final critical piece of information is the formula for choosing n, which determines the sizes of the matrices $\mathbf{M}_{n}^{\mathbf{0}}(\boldsymbol{\mu})$, $\mathbf{M}_{n-1}^{\mathbf{0}}(\boldsymbol{\mu})$, and $\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\boldsymbol{\mu})$.

Recall that in the construction of every SDP that appears in this paper, we are setting to zero the time derivatives of all moments up through order $m \in \mathbb{N}$. We specify this condition with the equation $\mathbf{A}\boldsymbol{\mu}_M = \mathbf{0}$, where $\boldsymbol{\mu}_M$ is the vector of all moments up through $M \equiv m + q - 1$, and q is the order of the highest-order reaction in the system. Inspecting the definition of $\mathbf{M}_n^{\mathbf{0}}(\boldsymbol{\mu})$ given in Section B.0.2, we see that the highest-order moments appearing in this matrix are order 2n. We want to ensure that the order M moments appear in this matrix, so that they are properly constrained. In other words, we want to ensure that $2n \geq M$. On the other hand, we don't want to choose n to be excessively large, because (as pointed out in Section B.0.3) this results in large matrices. Both goals are satisfied by setting $n \equiv \lceil \frac{M}{2} \rceil$.

With this final piece of information, we can compute the various measures of size for any of the SDPs described in the paper. For example, consider SDP (3.30) for calculating a bound on $\langle X_i \rangle_{ss}$ for the Michaelis-Menten system with m = 2. For this system, q = 2. It follows that M = 2 + 2 - 1 = 3 and that $n = \lfloor \frac{3}{2} \rfloor = 2$. As mentioned in the previous paragraph, the matrix $\mathbf{M}_n^0(\boldsymbol{\mu})$ includes moments up to order 2n = 4. Inspecting all other constructs appearing in SDP (3.30), we see that no higher-order moments appear. Thus, for this case, the highest-order moment that appears as a decision variable is order h = 4. The total number of decision variables is exactly the number of moments up to order 4, which is $d = \begin{pmatrix} \hat{N}+4\\ \hat{N} \end{pmatrix} = \begin{pmatrix} 2+4\\ \hat{N} \end{pmatrix} = 15$. The equation $\mathbf{A}\boldsymbol{\mu}_M = \mathbf{0}$ sets to zero the time derivatives of all moments up to order m. So, the number of equations implied by $\mathbf{A}\boldsymbol{\mu}_M = \mathbf{0}$ is $l = \begin{pmatrix} \hat{N}+2\\ \hat{N} \end{pmatrix} = \begin{pmatrix} 2+2\\ 2 \end{pmatrix} = 6$. Finally, the largest matrix appearing in an LMI is $\mathbf{M}_n^{\mathbf{0}}(\boldsymbol{\mu})$, with the number of rows and columns given by $L = \begin{pmatrix} \hat{N}+n\\ \hat{N} \end{pmatrix} = \begin{pmatrix} 2+2\\ 2 \end{pmatrix} = 6$. The size of the other SDPs appearing in the paper can be analyzed similarly.

B.0.6 Scaling

As noted in Section 3.3.5, one difficulty with using moments as decision variables in SDPs is that the moments of a distribution often differ by orders of magnitude. This can cause numerical difficulties for solvers which tend to treat all variables equally. Fortunately, these difficulties can usually be alleviated by appropriately *scaling* the decision variables. In this section we describe how to scale SDP (3.30). The procedure for scaling the other SDPs that appear in the paper is similar.

The problem and scaling solution is well-illustrated by a system with two independent molecular counts \hat{x}_1 and \hat{x}_2 . Suppose the steady-state probability distribution for the system $P_{\rm ss}(\cdot)$ is a Dirac distribution where all of the probability is concentrated on the single point $(\hat{x}_1, \hat{x}_2) = (3, 2)$. Then, the moments of the system are given by

$$\mu_{\mathbf{j}} = \sum_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}} \hat{\mathbf{x}}^{\mathbf{j}} P_{\mathrm{ss}}(\hat{\mathbf{x}}) = 3^{j_1} 2^{j_2}, \quad \forall \mathbf{j} = (j_1, j_2) \in \mathbb{N}^2.$$
(B.34)

The numerical values for these moments increase exponentially with the sub-indices j_1 and j_2 , leading to moments of different orders of magnitude.

However, suppose we define the scaled variables

$$x'_1 \equiv \frac{1}{3}\hat{x}_1, \quad x'_2 \equiv \frac{1}{2}\hat{x}_2.$$
 (B.35)

Then, for every point $\hat{\mathbf{x}}$ in the set of reachable states $\hat{\mathcal{X}}$, there is a corresponding

scaled state $\mathbf{x}' = \mathbf{D}^{-1} \hat{\mathbf{x}}$, where \mathbf{D} is the diagonal matrix

$$\mathbf{D} = \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} \equiv \begin{bmatrix} 3 \\ 2 \end{bmatrix}. \tag{B.36}$$

We will let \mathcal{X}' denote the set of these scaled reachable states. We can define a probability distribution over \mathcal{X}' :

$$P'_{\rm ss}(\mathbf{x}') \equiv P_{\rm ss}(\mathbf{D}\mathbf{x}'), \quad \forall \mathbf{x}' \in \mathcal{X}'.$$
 (B.37)

Furthermore, we can define the moments of this probability distribution,

$$\mu'_{\mathbf{j}} \equiv \sum_{\mathbf{x}' \in \mathcal{X}'} (\mathbf{x}')^{\mathbf{j}} P'_{\mathrm{ss}}(\mathbf{x}'), \quad \forall \mathbf{j} \in \mathbb{N}^2,$$
(B.38)

and we see that

$$\mu'_{\mathbf{j}} \equiv \sum_{\mathbf{x}' \in \mathcal{X}'} (\mathbf{x}')^{\mathbf{j}} P_{\mathrm{ss}}(\mathbf{D}\mathbf{x}'),$$

$$= \sum_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}} (\mathbf{D}^{-1} \hat{\mathbf{x}})^{\mathbf{j}} P_{\mathrm{ss}}(\hat{\mathbf{x}}),$$

$$= \left(\frac{1}{3}\right)^{j_1} \left(\frac{1}{2}\right)^{j_2} \sum_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}} \hat{\mathbf{x}}^{\mathbf{j}} P_{\mathrm{ss}}(\hat{\mathbf{x}}),$$

$$= \left(\frac{1}{3}\right)^{j_1} \left(\frac{1}{2}\right)^{j_2} \mu_{\mathbf{j}},$$

$$= \left(\frac{1}{3}\right)^{j_1} \left(\frac{1}{2}\right)^{j_2} 3^{j_1} 2^{j_2},$$

$$= 1,$$
(B.39)

for all multi-indices $\mathbf{j} \in \mathbb{N}^2$. That is, all the moments are equal to one. The fact that they are all equal is attributable to the special (Dirac) structure of original distribution $P_{\rm ss}(\cdot)$. More important than their equality is the fact that they are all of the same order of magnitude. This similarity in order of magnitude is what we will try to achieve for the moments appearing in SDP (3.30).

To apply the scaling idea of the previous example to SDP (3.30), we need to move to the more general case where there are \hat{N} independent species. Let $\mathbf{d} \equiv (d_1, \ldots, d_{\hat{N}})$ be the diagonal elements of the matrix $\mathbf{D} \in \mathbb{R}^{\hat{N} \times \hat{N}}$, where $d_i > 0$ is the scaling coefficient associated with independent species *i*. As before, we have a probability distribution $P'_{ss}(\cdot)$ over a set of scaled reachable states \mathcal{X}' and the corresponding scaled moments $\boldsymbol{\mu}'$. Applying the same logic as in Equation (B.39), we can conclude that the relationship between the scaled moments $\boldsymbol{\mu}'$ and the original moments $\boldsymbol{\mu}$ is given by

$$\mu_{\mathbf{j}} = \mathbf{d}^{\mathbf{j}} \mu_{\mathbf{j}}', \quad \forall \mathbf{j} \in \mathbb{N}^{\hat{N}}.$$
(B.40)

We now review each part of SDP (3.30) in turn, showing how each part can be expressed in terms of the scaled moments μ' . The objective function $\mu_{\mathbf{e}_i}$ is equal to $\mathbf{d}^{\mathbf{e}_i}\mu'_{\mathbf{e}_i} = d_i\mu'_{\mathbf{e}_i}$. So that the objective function is a reasonable order of magnitude, the factor of d_i is pulled outside of the optimization, as shown in SDP (B.42) below. To express the constraint $\mathbf{A}\mu_M = \mathbf{0}$ in terms of μ' , we need to multiply each column of the matrix \mathbf{A} by the appropriate scaling factor $\mathbf{d}^{\mathbf{j}}$, where \mathbf{j} is the multi-index corresponding to the current column. Doing so, we obtain the scaled matrix \mathbf{A}' and the equivalent equation $\mathbf{A}'\mu'_M = \mathbf{0}$. The equation $\mu_0 = 1$ translates simply into $\mu'_0 = 1$. All that remains is the LMIs. We can derive LMIs in terms of μ' in much the scaled states \mathbf{x}' and the probability distribution $P'_{ss}(\cdot)$ instead of $\hat{\mathbf{x}}$ and $P_{ss}(\cdot)$. This derivation results in the following LMIs:

$$\mathbf{M}_{n}^{\mathbf{0}}(\boldsymbol{\mu}') \succeq \mathbf{0},$$

$$\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\boldsymbol{\mu}') \succeq \mathbf{0}, \quad \forall j \in \{1, \dots, \hat{N}\},$$

$$\alpha_{k}\mathbf{M}_{n-1}^{\mathbf{0}}(\boldsymbol{\mu}') - \sum_{j=1}^{\hat{N}} \beta_{k,j} d_{j} \mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\boldsymbol{\mu}') \succeq \mathbf{0}, \quad \forall k \in \{1, \dots, L\}.$$
(B.41)

Essentially, the only change is that the moments μ' have been substituted in place of μ , and a factor of d_j has appeared in the third LMI. Bringing it all together, the resulting scaled SDP is

$$\frac{\langle x_i \rangle_{ss}^U}{d_i} = \max_{\boldsymbol{\mu}'} \quad \boldsymbol{\mu}'_{\mathbf{e}_i}
\text{s.t.} \quad \mathbf{A}' \boldsymbol{\mu}'_M = \mathbf{0},
\qquad \boldsymbol{\mu}'_0 = 1,
\qquad \mathbf{M}_n^{\mathbf{0}}(\boldsymbol{\mu}') \succeq \mathbf{0},
\qquad \mathbf{M}_{n-1}^{\mathbf{e}_j}(\boldsymbol{\mu}') \succeq \mathbf{0}, \quad \forall j \in \{1, \dots, \hat{N}\},
\qquad \boldsymbol{\alpha}_k \mathbf{M}_{n-1}^{\mathbf{0}}(\boldsymbol{\mu}') - \sum_{j=1}^{\hat{N}} \beta_{k,j} d_j \mathbf{M}_{n-1}^{\mathbf{e}_j}(\boldsymbol{\mu}') \succeq \mathbf{0}, \quad \forall k \in \{1, \dots, L\}.$$
(B.42)

We have shown how to incorporate scaling factors **d** into SDP (3.30). What we haven't explained is how to pick the scaling factors appropriately. There is no strict rule. However, as a heuristic we suggest that the scaling factors be picked so that $d_i \approx \langle x_i \rangle_{ss}$. Granted, we usually won't know $\langle x_i \rangle_{ss}$ for each independent species *i* a priori (if we did, we wouldn't need to solve SDP (3.30)), but an order of magnitude estimate will suffice. This estimate could come from a variety of sources – for example, solution of the deterministic rate law equations, experimental data, prior literature, or a just a few kinetic monte carlo simulations. In our experience, using an overestimate of $\langle x_i \rangle_{ss}$ for the scaling factor d_i usually leads to more numerically stable SDPs than using an under-estimate. Accordingly, for most of the examples in the paper, for the scaling factor d_i , we used a crude over-estimate of $\langle x_i \rangle_{ss}$ implied by the stoichiometry and initial condition, obtained through solving the following linear program:

$$d_{i} = \max_{\hat{\mathbf{x}}} \quad \hat{x}_{i}$$

s.t. $\hat{\mathbf{x}} \ge \mathbf{0},$ (B.43)
 $\tilde{\mathbf{x}}_{0} + \tilde{\mathbf{B}}^{-1}\hat{\mathbf{B}}(\hat{\mathbf{x}}_{0} - \hat{\mathbf{x}}) \ge \mathbf{0}$

The two constraints simply specify that the molecular counts of the independent and dependent species must be nonnegative. (See Equation (4.5) and the surrounding discussion.) Of course, this method only works when the molecular count of species i is bounded above. This is guaranteed to be true only for reaction systems that are closed with respect to mass. For open systems, we could employ a state-space analysis as proposed by Gupta et al [29] to determine which species are bounded and which are "free". We could then apply the above LP to the bounded species. For the free species, we must rely on one of the aforementioned alternative methods for estimating $\langle x_i \rangle_{ss}$.

B.0.7 LMIs to Exclude Non-integer States

In Section IV B, we stated that we could selectively add LMIs reflecting the integrality requirement and thus obtain tighter bounds. We now demonstrate this idea for the Reversible Dimerization system of Section IV B. Recall that there was a gap in the bounds on $\langle A \rangle_{ss}$ (see Figure 3-2, top plot). In particular, for $c_2 = 0$, we had $\langle A \rangle_{ss}^L = 0$ and $\langle A \rangle_{ss}^U = 1$. We claimed that this gap was attributable to a nonphysical distribution, where all of the probability is concentrated on the point $x_A = 1$, implying 4.5 molecules of B. The moments of this distribution $\boldsymbol{\mu}^* = (1, 1, 1, ...)$ are feasible for SDP (3.30), because they satisfy the stationarity condition $A\boldsymbol{\mu}^* = \mathbf{0}$, and they satisfy the necessary conditions for the moments of any distribution supported on the polyhedral set $\bar{\mathcal{X}}$, which for this reaction system reduces to

$$\bar{\mathcal{X}} = \left\{ x_{\mathrm{A}} \in \mathbb{R} : \ 0 \le x_{\mathrm{A}} \le 10 \right\}$$
(B.44)

(see Section 3.3.4). However, from an analysis of the stoichiometry of this system, we know that the state $x_{\rm A} = 1$ is actually not possible. Furthermore, it is not possible that $x_{\rm A}$ takes any value in the open interval (0, 2). In other words, the steady-state probability distribution $P_{\rm ss}(\cdot)$ has no support in the interval (0, 2). This observation implies an additional necessary condition on the true steady-state moments μ :

$$\mathbf{M}_{n-2}^{2\mathbf{e}_1}(\boldsymbol{\mu}) - 2\mathbf{M}_{n-2}^{\mathbf{e}_1}(\boldsymbol{\mu}) \succeq \mathbf{0}, \tag{B.45}$$

for $n \ge 2$. We can add this LMI as a constraint to SDP (3.30), further restricting the feasible set, and, in particular, excluding the moments μ^* .

The derivation of LMI (B.45) follows the same pattern as the derivation of the LMIs in Section B.0.3, though the notation is simplified somewhat because $\hat{N} = 1$. First, we construct an arbitrary polynomial

$$p(x_{\rm A}) \equiv \sum_{j=0}^{n-2} p_j \hat{\mathbf{x}}^j. \tag{B.46}$$

Then, because $P_{\rm ss}(\cdot) = 0$ in the open interval (0, 2), we have

$$\sum_{x_{\rm A}\in\hat{\mathcal{X}}} x_{\rm A}(x_{\rm A}-2)p^2(x_{\rm A})P_{\rm ss}(x_{\rm A}) \ge 0.$$
(B.47)

Expanding the squared polynomial and using the definitions of $\mathbf{M}_{n-2}^{2\mathbf{e}_1}(\boldsymbol{\mu})$ and $\mathbf{M}_{n-2}^{\mathbf{e}_1}(\boldsymbol{\mu})$, this becomes

$$\mathbf{p}^{\mathrm{T}}\left(\mathbf{M}_{n-2}^{2\mathbf{e}_{1}}(\boldsymbol{\mu}) - 2\mathbf{M}_{n-2}^{\mathbf{e}_{1}}(\boldsymbol{\mu})\right)\mathbf{p} \ge 0, \tag{B.48}$$

where **p** is a vector of the coefficients of the polynomial $p(\cdot)$. The fact that these coefficients were arbitrary implies LMI (B.45).

Adding LMI (B.45) to SDP (3.30) we obtain

$$\langle x_i \rangle_{ss}^U = \max_{\boldsymbol{\mu}} \quad \boldsymbol{\mu}_{\mathbf{e}_i}$$
s.t. $\mathbf{A} \boldsymbol{\mu}_M = \mathbf{0},$

$$\boldsymbol{\mu}_{\mathbf{0}} = 1,$$

$$\mathbf{M}_n^{\mathbf{0}}(\boldsymbol{\mu}) \succeq \mathbf{0}, \qquad (B.49)$$

$$\mathbf{M}_{n-1}^{\mathbf{e}_j}(\boldsymbol{\mu}) \succeq \mathbf{0}, \quad \forall j \in \{1, \dots, \hat{N}\},$$

$$\alpha_k \mathbf{M}_{n-1}^{\mathbf{0}}(\boldsymbol{\mu}) - \sum_{j=1}^{\hat{N}} \beta_{k,j} \mathbf{M}_{n-1}^{\mathbf{e}_j}(\boldsymbol{\mu}) \succeq \mathbf{0}, \quad \forall k \in \{1, \dots, L\},$$

$$\mathbf{M}_{n-2}^{2\mathbf{e}_1}(\boldsymbol{\mu}) - 2\mathbf{M}_{n-2}^{\mathbf{e}_1}(\boldsymbol{\mu}) \succeq \mathbf{0}.$$

Repeatedly solving this SDP with m = 2 as we did in Section IV B, we obtain the bottom plot of Figure 3-2. Comparing the top and bottom plots of Figure 3-2, we

| Rate Constant | Value (s^{-1}) |
|------------------------|-------------------|
| c_1 | 1 |
| c_2 | 1 |
| c_3 | 1 |
| c_4 | 1 |
| c_5 | 1 |
| c_6 | 1 |
| c_7 | 1×10^4 |
| C ₈ | 1×10^{4} |
| C_9 | 1 |
| c_{10} | 1 |
| c_{11} | 1 |
| c_{12} | 1×10^{5} |
| c_{13} | 1×10^{5} |
| <i>c</i> ₁₄ | 1 |

Table B.2: Reaction rate constants for the Larger Example

see that, for low values of c_2 , the upper bound $\langle A \rangle_{ss}^U$ is tighter in the bottom plot, converging to zero as $c_2 \to 0$.

B.0.8 Rate Constants for the Larger Example

The larger reaction system example shown in Figure 3-4 can also be represented as

$$A \xrightarrow{c_{1}} B + C \qquad A \xrightarrow{c_{9}} H$$

$$B \xrightarrow{c_{3}} 2D \qquad H \xrightarrow{c_{11}} I$$

$$C \xrightarrow{c_{5}} E \qquad I \xrightarrow{c_{12}} J$$

$$E + F \xrightarrow{c_{7}} G \qquad J \xrightarrow{c_{14}} A$$
(B.50)

The rate constants for these reactions are given in Table B.2.

B.0.9 Relative Probabilities for the Schlögl System

In Section 3.8.3, we plotted relative probabilities for the Schlögl System (3.34). We now explain how these relative probabilities were obtained.

As explained in Section 3.8.3, the counts of species A and B are assumed constant, so the only variable is the molecular count of species X. Accordingly, we will denote the molecular counts of species A and B using the symbols A and B, respectively, and we will denote the molecular count of X as simply x. Futhermore, for notational simplicity, we will denote the steady-state probability distribution $P_{ss}(\cdot)$ as $P(\cdot)$. The CME (3.6) for this system at steady state is

$$0 = P(x - 1)a_1(x - 1) - P(x)a_1(x) + P(x + 1)a_2(x + 1) - P(x)a_2(x) + P(x - 1)a_3(x - 1) - P(x)a_3(x) + P(x + 1)a_4(x + 1) - P(x)a_4(x), \quad \forall x \in \mathbb{N}.$$
(B.51)

where each row of this equation corresponds to one reaction. Rearranging to collect the coefficients associated with P(x - 1), P(x), and P(x + 1), we obtain

$$0 = (a_1(x-1) + a_3(x-1)) P(x-1)$$

- $(a_1(x) + a_2(x) + a_3(x) + a_4(x)) P(x)$
+ $(a_2(x+1) + a_4(x+1)) P(x+1), \quad \forall x \in \mathbb{N}.$ (B.52)

The propensity functions for the four reactions of the system are

• Reaction 1:

$$a_1(x) = c_1 {\binom{x}{2}} {\binom{A}{1}} = c_1 A {\binom{x}{2}}.$$
 (B.53)

• Reaction 2:

$$a_2(x) = c_2 \binom{x}{3}.\tag{B.54}$$

• Reaction 3:

$$a_3(x) = c_3 \binom{B}{1} = c_3 B.$$
 (B.55)

• Reaction 4:

$$a_4(x) = c_4 \begin{pmatrix} x \\ 1 \end{pmatrix} = c_4 x.$$
 (B.56)

Substituting these propensity functions into the CME, we obtain

$$0 = \left(c_1 A \begin{pmatrix} x - 1 \\ 2 \end{pmatrix} + c_3 B\right) P(x - 1) - \left(c_1 A \begin{pmatrix} x \\ 2 \end{pmatrix} + c_2 \begin{pmatrix} x \\ 3 \end{pmatrix} + c_3 B + c_4 x\right) P(x) + \left(c_2 \begin{pmatrix} x + 1 \\ 3 \end{pmatrix} + c_4 (x + 1)\right) P(x + 1), \quad \forall x \in \mathbb{N},$$
(B.57)

where we adopt the convention that

$$\binom{i}{j} = 0 \tag{B.58}$$

if i < j.

Now, we assume (reasonably) that P(-1) = 0. That is, there will never be -1 molecules of X. Then, we can simplify Equation (B.57) for the special case of x = 0 to obtain

$$0 = -c_3 BP(0) + c_4 P(1). \tag{B.59}$$

If we fix the value of P(0) to some arbitrary (nonnegative) value, the above equation fixes the value of P(1), For the special case of x = 1, Equation (B.57) gives

$$0 = c_3 BP(0) - (c_3 B + c_4) P(1) + 2c_4 P(2).$$
(B.60)

Since the values of P(0) and P(1) have been fixed, this equation fixes the value of P(2). We can then use Equation (B.57) to show that P(1) and P(2) fix the value of P(3), and so on. The conclusion is this: once we fix a value of P(0), the entire sequence $\{P(x)\}_{x=1}^{\infty}$ is fixed by Equation (B.57). Moreover, because of the linearity of Equation (B.57), every element of the sequence $\{P(x)\}_{x=1}^{\infty}$ is proportional to p_0 . This implies that $P(0) \neq 0$, because P(0) = 0 would imply that P(x) = 0 for all $x \in \mathbb{N}$,

and $\sum_{x \in \mathbb{N}} P(x) = 0 \neq 1$. So P(0) > 0. This means that the relative probability ratio

$$\frac{P(x)}{P(0)} \equiv \tilde{P}(x) \tag{B.61}$$

is well-defined for each $x \in \mathbb{N}$. Furthermore, if we divide Equation (B.57) by P(0),

$$0 = \left(c_1 A \begin{pmatrix} x - 1 \\ 2 \end{pmatrix} + c_3 B\right) \tilde{P}(x - 1)$$

$$- \left(c_1 A \begin{pmatrix} x \\ 2 \end{pmatrix} + c_2 \begin{pmatrix} x \\ 3 \end{pmatrix} + c_3 B + c_4 x\right) \tilde{P}(x)$$

$$+ \left(c_2 \begin{pmatrix} x + 1 \\ 3 \end{pmatrix} + c_4 (x + 1)\right) \tilde{P}(x + 1), \quad \forall x \in \mathbb{N}.$$
 (B.62)

we obtain a system of equations that fixes the relative probability ratio $\tilde{P}(x)$ for every $x \in \mathbb{N}$. To see this, we just have to consider the equations for $x = 0, 1, 2, \ldots$, in turn, much like we did before, using the fact that $\tilde{P}(0) = 1$. Doing so, we can calculate the relative probability ratios $\tilde{P}(1), \ldots, \tilde{P}(N)$ for an arbitrarily large $N \in \mathbb{N}$. These probability ratios can give us an accurate picture of the shape of the steady-state probability distribution over the set $x_{\rm X} = 0, 1, \ldots, N$. In particular, we can see how this distribution transitions from unimodal to bimodal and back to unimodal as c_4 varies from 3 s⁻¹ to 4 s⁻¹. See Figure B-1.



Figure B-1: Relative probabilities for the Schlögl System over a range of parameter values for c_4 . The values of the other parameters are as stated in Section3.4.3.

B.0.10 Number of States in the Larger Example

In Section 3.4.4, we claimed that the number of states accessible to larger reaction system shown in Figure 3-4 is 74,816,108,146. In this section, we explain how we arrived at this number.

The calculation is easier to understand if we represent the reaction system as shown in Figure B-2, which divides the reaction system into three subsystems. Recall that we specified an initial condition of 100 molecules of A and 100 molecules of F.



Figure B-2: Alternative representation of Figure 3-4

First, from Figure B-2, we see that the molecules of F essentially just enable the conversion of E to G. Furthermore, since this reaction requires 1 molecule of F for every molecule of E converted, and since we can have at most 100 molecules of E, we can effectively forget about the F molecules, and treat the conversion as simply $E \rightleftharpoons G$. The count of F molecules at any time is simply $x_F = 100 - x_G$.

We can now focus on the 100 initial A molecules and consider how they might distribute themselves among Subsystems 1, 2, and 3. Suppose that $n \in \{0, ..., 100\}$ molecules of A have been converted to B and C molecules. Then, the problem of counting the number of states accessible to the whole system reduces to three simpler sub-problems:

- How many states are accessible to Subsystem 1 given an initial count of 100 n molecules of A?
- How many states are accessible to Subsystem 2 given an initial count of *n* molecules of B?
- How many states are accessible to Subsystem 3 given an initial count of *n* molecules of C?

To answer the first question, we just have to consider how many ways the 100 - n molecules can distribute themselves among the four different molecular "bins", A, H, I, and J. Using the classic "stars and bars" argument from combinatorics, we see that the number of states accessible to Subsystem 1 is $\binom{100-n+4-1}{4-1} = \binom{100-n+3}{3}$. By the same argument, we can conclude that there are $\binom{n+1}{1}$ states accessible to Subsystem 2 and $\binom{n+2}{2}$ states accessible to Subsystem 3. For a given n, the state of each subsystem is independent of the state of the other two. So, given n, the number of states accessible to the whole reaction system is the product of these binomial coefficients: $\binom{100-n+3}{3}\binom{n+1}{1}\binom{n+2}{2}$. To calculate the total number of states, we just have to sum over all possible values of n:

number of states
$$=\sum_{n=0}^{100} \binom{100-n+3}{3} \binom{n+1}{1} \binom{n+2}{2}$$
 (B.63)

We computed this sum using a simple "for" loop in Matlab, obtaining 74,816,108,146.

B.0.11 Derivation of SDP (3.39)

In this section, we derive the SDP for the upper bound on the steady-state probability that $\hat{X}_i \in [x_{\min}, x_{\max}]$ for independent species $i \in \{1, \ldots, \hat{N}\}$. In this derivation, the set of reachable states $\hat{\mathbf{x}} \in \hat{\mathcal{X}}$ such that \hat{x}_i lies in the interval $[x_{\min}, x_{\max}]$ plays a special a role. We will need to refer to this set often. So, for notational convenience, we identify it with the following symbol:

$$\hat{\mathcal{X}}_i \equiv \{ \hat{\mathbf{x}} \in \hat{\mathcal{X}} : \hat{x}_i \in [x_{\min}, x_{\max}] \}.$$
(B.64)

With this definition, we can construct the following optimization problem for calculating an upper bound on the probability that $\hat{X}_i \in [x_{\min}, x_{\max}]$:

$$\begin{aligned}
\max_{P_{\rm ss}} &= \sum_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}_i} P_{\rm ss}(\hat{\mathbf{x}}) \\
\text{s.t.} &\sum_{r=1}^R \left[P_{\rm ss}(\hat{\mathbf{x}} - \hat{\mathbf{s}}_r) a_r(\hat{\mathbf{x}} - \hat{\mathbf{s}}_r) - P_{\rm ss}(\hat{\mathbf{x}}) a_r(\hat{\mathbf{x}}) \right] = 0, \quad \forall \hat{\mathbf{x}} \in \hat{\mathcal{X}}, \\
&\sum_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}} P_{\rm ss}(\hat{\mathbf{x}}) = 1, \\
&P_{\rm ss} : \hat{\mathcal{X}} \to \mathbb{R}_+.
\end{aligned}$$
(B.65)

The variable in this optimization is a function $P_{\rm ss}$ which maps each reachable state $\hat{\mathbf{x}} \in \hat{\mathcal{X}}$ to a nonnegative number. The constraint $\sum_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}} P_{\rm ss}(\hat{\mathbf{x}}) = 1$ ensures that these nonnegative numbers must sum to one, and thus allows them to be interpreted as probabilities. The first constraint is simply the CME (3.6) in which the time derivatives for all states $\hat{\mathbf{x}} \in \hat{\mathcal{X}}$ have been set to zero, specifying that $P_{\rm ss}(\cdot)$ is a steady-state distribution. The objective function is a sum over only those states $\hat{\mathbf{x}} \in \hat{\mathcal{X}}$ that satisfy $\hat{x}_i \in [x_{\min}, x_{\max}]$, giving the total probability associated with the probability that $\hat{X}_i \in [x_{\min}, x_{\max}]$.

The next step is to decompose $P_{ss}(\cdot)$ into two distributions $Q : \hat{\mathcal{X}} \to \mathbb{R}_+$ and $S : \hat{\mathcal{X}} \to \mathbb{R}_+$ such that

$$P_{\rm ss}(\hat{\mathbf{x}}) = Q(\hat{\mathbf{x}}) + S(\hat{\mathbf{x}}), \quad \forall \hat{\mathbf{x}} \in \hat{\mathcal{X}}$$
(B.66)

and

$$S(\hat{\mathbf{x}}) = 0, \quad \forall \hat{\mathbf{x}} \notin \hat{\mathcal{X}}_i.$$
 (B.67)

This latter equation implies that the distribution S is supported only on $\hat{\mathcal{X}}_i$. The reason for this decomposition may not be clear yet, but will be later. Rewriting Problem (B.65) in terms of the two distributions $Q(\cdot)$ and $S(\cdot)$, we obtain

$$\max_{Q,S} = \sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}_{i}} Q(\hat{\mathbf{x}}) + \sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}_{i}} S(\hat{\mathbf{x}})$$
s.t.
$$\sum_{r=1}^{R} [Q(\hat{\mathbf{x}} - \hat{\mathbf{s}}_{r})a_{r}(\hat{\mathbf{x}} - \hat{\mathbf{s}}_{r}) - Q(\hat{\mathbf{x}})a_{r}(\hat{\mathbf{x}})]$$

$$+ \sum_{r=1}^{R} [S(\hat{\mathbf{x}} - \hat{\mathbf{s}}_{r})a_{r}(\hat{\mathbf{x}} - \hat{\mathbf{s}}_{r}) - S(\hat{\mathbf{x}})a_{r}(\hat{\mathbf{x}})] = 0, \quad \forall \hat{\mathbf{x}} \in \hat{\mathcal{X}}, \quad (B.68)$$

$$\sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}} Q(\hat{\mathbf{x}}) + \sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}} S(\hat{\mathbf{x}}) = 1,$$

$$Q: \hat{\mathcal{X}} \to \mathbb{R}_{+}, \quad S: \hat{\mathcal{X}} \to \mathbb{R}_{+},$$

$$S(\hat{\mathbf{x}}) = 0, \quad \forall \hat{\mathbf{x}} \notin \hat{\mathcal{X}}_{i}.$$

Problems (B.65) and (B.68) are equivalent in the sense that starting from a feasible solution of either, we can construct a feasible solution for the other with the same objective value. For example, if Q and S are feasible for Problem (B.68), then $P_{ss} \equiv$ Q + S is feasible for Problem (B.65) and has the same objective value. Going the other way, if P_{ss} is feasible for Problem (B.65), then $Q \equiv P_{ss}$ and $S \equiv 0$ is feasible for Problem (B.68). It follows that the two problems have the same optimal objective value.

Taking this idea one step further, Problem (B.68) is equivalent to the following optimization problem, which differs from Problem (B.68) only in the objective function:

$$\max_{Q,S} = \sum_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}} S(\hat{\mathbf{x}})$$
s.t.
$$\sum_{r=1}^{R} [Q(\hat{\mathbf{x}} - \hat{\mathbf{s}}_{r})a_{r}(\hat{\mathbf{x}} - \hat{\mathbf{s}}_{r}) - Q(\hat{\mathbf{x}})a_{r}(\hat{\mathbf{x}})]$$

$$+ \sum_{r=1}^{R} [S(\hat{\mathbf{x}} - \hat{\mathbf{s}}_{r})a_{r}(\hat{\mathbf{x}} - \hat{\mathbf{s}}_{r}) - S(\hat{\mathbf{x}})a_{r}(\hat{\mathbf{x}})] = 0, \quad \forall \hat{\mathbf{x}} \in \hat{\mathcal{X}}, \quad (B.69)$$

$$\sum_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}} Q(\hat{\mathbf{x}}) + \sum_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}} S(\hat{\mathbf{x}}) = 1,$$

$$Q : \hat{\mathcal{X}} \to \mathbb{R}_{+}, \quad S : \hat{\mathcal{X}} \to \mathbb{R}_{+},$$

$$S(\hat{\mathbf{x}}) = 0, \quad \forall \hat{\mathbf{x}} \notin \hat{\mathcal{X}}_{i}.$$

To see this equivalence, suppose that Q and S are feasible for Problem (B.69). Without loss of generality, we may assume that $Q(\hat{\mathbf{x}}) = 0$ for all $\hat{\mathbf{x}} \in \hat{\mathcal{X}}_i$. Otherwise, we could construct another feasible solution $\tilde{Q} : \hat{\mathcal{X}} \to \mathbb{R}_+$ and $\tilde{S} : \hat{\mathcal{X}} \to \mathbb{R}_+$ which gives a greater objective value. In particular, we would let

$$\tilde{Q}(\hat{\mathbf{x}}) \equiv \begin{cases} 0, & \text{if } \hat{\mathbf{x}} \in \hat{\mathcal{X}}_i \\ Q(\hat{\mathbf{x}}), & \text{otherwise.} \end{cases}$$
(B.70)

$$\tilde{S}(\hat{\mathbf{x}}) \equiv \begin{cases} Q(\hat{\mathbf{x}}) + S(\hat{\mathbf{x}}), & \text{if } \hat{\mathbf{x}} \in \hat{\mathcal{X}}_i \\ 0, & \text{otherwise.} \end{cases}$$
(B.71)

Continuing with the assumption that $Q(\hat{\mathbf{x}}) = 0$ for all $\hat{\mathbf{x}} \in \hat{\mathcal{X}}_i$, and noting that feasibility implies $S(\hat{\mathbf{x}}) = 0$ for all $\hat{\mathbf{x}} \notin \hat{\mathcal{X}}_i$, we see that the two objective functions are equal:

$$\sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}_i} Q(\hat{\mathbf{x}}) + \sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}_i} S(\hat{\mathbf{x}}) = \sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}_i} S(\hat{\mathbf{x}}) = \sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}} S(\hat{\mathbf{x}}).$$
(B.72)

This argument shows that the optimal value of Problem (B.68) is greater than or equal to the optimal value of Problem (B.69). To show the reverse inequality, suppose that Q and S are feasible for Problem (B.68). Then, the functions $\tilde{Q}(\cdot)$ and $\tilde{S}(\cdot)$ defined above are feasible for Problem (B.69) and have the same objective value. This establishes the reverse inequality. The two inequalities together imply that the optimal values of Problems (B.68) and (B.69) are equal. Furthermore, since Problem (B.68) is equivalent to our original problem, Problem (B.65), the optimal value of Problem (B.69) is also an upper bound on the probability that $\hat{X}_i \in [x_{\min}, x_{\max}]$.

The next step is to recast Problem (B.69) in terms of the moments of the two distributions. Let the moments of the distribution $Q(\cdot)$ and $S(\cdot)$ be given by

$$y_{\mathbf{j}} \equiv \sum_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}} \hat{\mathbf{x}}^{\mathbf{j}} Q(\hat{\mathbf{x}}), \quad z_{\mathbf{j}} \equiv \sum_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}} \hat{\mathbf{x}}^{\mathbf{j}} S(\hat{\mathbf{x}}), \quad \forall \mathbf{j} \in \mathbb{N}^{\hat{N}}.$$
 (B.73)

Then, we can rewrite Problem (B.69) as

$$\begin{aligned}
\max_{Q,S,\mathbf{y},\mathbf{z}} &= z_{\mathbf{0}} \\
\text{s.t.} \quad \mathbf{A}\mathbf{y}_{M} + \mathbf{A}\mathbf{z}_{M} &= \mathbf{0}, \\
& y_{\mathbf{0}} + z_{\mathbf{0}} &= 1, \\
& Q : \hat{\mathcal{X}} \to \mathbb{R}_{+}, \quad S : \hat{\mathcal{X}} \to \mathbb{R}_{+}, \\
& S(\hat{\mathbf{x}}) &= 0, \quad \forall \hat{\mathbf{x}} \notin \hat{\mathcal{X}}_{i}, \\
& y_{\mathbf{j}} &\equiv \sum_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}} \hat{\mathbf{x}}^{\mathbf{j}} Q(\hat{\mathbf{x}}), \quad z_{\mathbf{j}} &\equiv \sum_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}} \hat{\mathbf{x}}^{\mathbf{j}} S(\hat{\mathbf{x}}), \quad \forall \mathbf{j} \in \mathbb{N}^{\hat{N}}.
\end{aligned}$$
(B.74)

Our list of decision variables now includes the functions $Q(\cdot)$, $S(\cdot)$ and their respective moment sequences \mathbf{y}, \mathbf{z} . The last line of constraints enforces the relationships between theses variables. The two preceding lines of constraints are exactly the same as those appearing in Problem (B.69). The second line of constraints enforces that the total probability is one, and, given our definition of the moments, is equivalent to the second line of constraints appearing in Problem (B.69). The first line of constraints reflects the steady state conditions of Problem (B.69). However, the two are not equivalent. While the equation $\mathbf{Ay}_M + \mathbf{Az}_M = \mathbf{0}$ species that the time derivatives of all moments up through order $m \in \mathbb{N}$ are equal to zero, the steady-state conditions in the first line of constraints in Problem (B.69) imply that the time derivatives of *all* moments are equal to zero. Thus, $\mathbf{Ay}_M + \mathbf{Az}_M = \mathbf{0}$ is a relaxation of the steady-state conditions, with the extent of the relaxation depending on our choice of m. Lastly, the objective
functions of Problems (B.74) and (B.69) are equivalent. Because of the relaxation of the steady-state conditions, the optimal values of Problems (B.74) and (B.69) are not necessarily equal. However, the former is guaranteed to be an upper bound on the latter. This means that the optimal value of Problem (B.74) is an upper bound on the probability that $\hat{X}_i \in [x_{\min}, x_{\max}]$, just perhaps not the tightest upper bound.

The second-to-last step in the derivation of SDP (3.39) is to add necessary conditions on the moment sequences \mathbf{y} and \mathbf{z} as constraints. The first set of these constraints are the LMIs (3.27) - (3.29), written for both \mathbf{y} and \mathbf{z} :

$$\mathbf{M}_{n}^{\mathbf{0}}(\mathbf{y}) \succeq \mathbf{0}, \quad \mathbf{M}_{n-1}^{\mathbf{0}}(\mathbf{z}) \succeq \mathbf{0}, \\
\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\mathbf{y}) \succeq \mathbf{0}, \quad \mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\mathbf{z}) \succeq \mathbf{0}, \quad \forall j \in \{1, \dots, \hat{N}\}, \\
\alpha_{k}\mathbf{M}_{n-1}^{\mathbf{0}}(\mathbf{y}) - \sum_{j=1}^{\hat{N}} \beta_{k,j}\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\mathbf{y}) \succeq \mathbf{0}, \quad \forall k \in \{1, \dots, L\}, \\
\alpha_{k}\mathbf{M}_{n-1}^{\mathbf{0}}(\mathbf{z}) - \sum_{j=1}^{\hat{N}} \beta_{k,j}\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\mathbf{z}) \succeq \mathbf{0}, \quad \forall k \in \{1, \dots, L\}.$$
(B.75)

As explained in Section (3.3.4), these are necessary conditions for the moments of any distribution supported on the polyhedral set $\hat{\mathcal{X}} \supset \hat{\mathcal{X}}$. The other necessary condition is also an LMI and reflects the fact that $S(\hat{\mathbf{x}}) = 0$ for all $\hat{\mathbf{x}} \notin \hat{\mathcal{X}}_i = \{\hat{\mathbf{x}} \in \hat{\mathcal{X}} : \hat{x}_i \in [x_{\min}, x_{\max}]\}$:

$$-\mathbf{M}_{n-2}^{2\mathbf{e}_{i}}(\mathbf{z}) + (x_{\max} + x_{\min})\mathbf{M}_{n-2}^{\mathbf{e}_{i}}(\mathbf{z}) - x_{\min}x_{\max}\mathbf{M}_{n-2}^{\mathbf{0}}(\mathbf{z}) \succeq \mathbf{0}.$$
 (B.76)

The derivation of this LMI closely parallels those of LMI (B.45) in Section B.0.7 and LMIs (3.27) - (3.29) in Section B.0.3. The key fact is that the polynomial $-(\hat{x}_i - x_{\min})(\hat{x}_i - x_{\max})$ is nonnegative for all $\hat{\mathbf{x}} \in \hat{\mathcal{X}}_i$ and negative for all $\hat{\mathbf{x}}$ such that $\hat{x}_i \notin [x_{\min}, x_{\max}]$. Adding these LMIs to Problem (B.74), we obtain

$$\begin{split} \max_{Q,S,\mathbf{y},\mathbf{z}} &= z_{\mathbf{0}} \\ \text{s.t.} & \mathbf{A}\mathbf{y}_{M} + \mathbf{A}\mathbf{z}_{M} = \mathbf{0}, \\ & y_{\mathbf{0}} + z_{\mathbf{0}} = 1, \\ & Q: \hat{\mathcal{X}} \to \mathbb{R}_{+}, \quad S: \hat{\mathcal{X}} \to \mathbb{R}_{+}, \\ & S(\hat{\mathbf{x}}) = 0, \quad \forall \hat{\mathbf{x}} \notin \hat{\mathcal{X}}_{i}, \\ & y_{\mathbf{j}} \equiv \sum_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}} \hat{\mathbf{x}}^{\mathbf{j}} Q(\hat{\mathbf{x}}), \quad z_{\mathbf{j}} \equiv \sum_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}} \hat{\mathbf{x}}^{\mathbf{j}} S(\hat{\mathbf{x}}), \quad \forall \mathbf{j} \in \mathbb{N}^{\hat{N}}, \\ & \mathbf{M}_{n}^{0}(\mathbf{y}) \succeq \mathbf{0}, \quad \mathbf{M}_{n-1}^{0}(\mathbf{z}) \succeq \mathbf{0} \\ & \mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\mathbf{y}) \succeq \mathbf{0}, \quad \mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\mathbf{z}) \succeq \mathbf{0}, \quad \forall j \in \{1, \dots, \hat{N}\}, \\ & \alpha_{k} \mathbf{M}_{n-1}^{\mathbf{0}}(\mathbf{y}) - \sum_{j=1}^{\hat{N}} \beta_{k,j} \mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\mathbf{y}) \succeq \mathbf{0}, \quad \forall k \in \{1, \dots, L\}, \\ & \alpha_{k} \mathbf{M}_{n-1}^{\mathbf{0}}(\mathbf{z}) - \sum_{j=1}^{\hat{N}} \beta_{k,j} \mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\mathbf{z}) \succeq \mathbf{0}, \quad \forall k \in \{1, \dots, L\}, \\ & - \mathbf{M}_{n-2}^{2\mathbf{e}_{i}}(\mathbf{z}) + (x_{\max} + x_{\min}) \mathbf{M}_{n-2}^{\mathbf{e}_{i}}(\mathbf{z}) - x_{\min} x_{\max} \mathbf{M}_{n-2}^{\mathbf{0}}(\mathbf{z}) \succeq \mathbf{0}. \end{split}$$

Since all of the LMIs we have added are necessary conditions for the moment sequences \mathbf{y} and \mathbf{z} , they do not change the feasible set. This means that the optimal value of Problem (B.77) is exactly the same as that of Problem (B.74).

The final step in the derivation of SDP (3.39) is to remove all mentions of the

distributions $Q(\cdot)$ and $S(\cdot)$ from the optimization problem

$$\begin{array}{ll}
\max_{\mathbf{y},\mathbf{z}} & z_{\mathbf{0}} \\
\text{s.t.} & \mathbf{A}\mathbf{y}_{M} + \mathbf{A}\mathbf{z}_{M} = \mathbf{0}, \quad y_{\mathbf{0}} + z_{\mathbf{0}} = 1, \\
& \mathbf{M}_{n}^{\mathbf{0}}(\mathbf{y}) \succeq \mathbf{0}, \quad \mathbf{M}_{n}^{\mathbf{0}}(\mathbf{z}) \succeq \mathbf{0} \\
& \mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\mathbf{y}) \succeq \mathbf{0}, \quad \mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\mathbf{z}) \succeq \mathbf{0}, \quad \forall j \in \{1, \dots, \hat{N}\}, \\
& \alpha_{k}\mathbf{M}_{n-1}^{\mathbf{0}}(\mathbf{y}) - \sum_{j=1}^{\hat{N}} \beta_{k,j}\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\mathbf{y}) \succeq \mathbf{0}, \quad \forall k \in \{1, \dots, L\}, \\
& \alpha_{k}\mathbf{M}_{n-1}^{\mathbf{0}}(\mathbf{z}) - \sum_{j=1}^{\hat{N}} \beta_{k,j}\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\mathbf{z}) \succeq \mathbf{0}, \quad \forall k \in \{1, \dots, L\}, \\
& - \mathbf{M}_{n-2}^{2\mathbf{e}_{i}}(\mathbf{z}) + (x_{\max} + x_{\min})\mathbf{M}_{n-2}^{\mathbf{e}_{i}}(\mathbf{z}) - x_{\min}x_{\max}\mathbf{M}_{n-2}^{\mathbf{0}}(\mathbf{z}) \succeq \mathbf{0}.
\end{array}$$
(3.39)

This is another relaxation, because the sequences \mathbf{y} and \mathbf{z} are no longer necessarily moments of some distribution. This means that the optimal value of SDP (3.39) is an upper bound on the optimal value of Problem (B.77). Since the optimal value of Problem (B.77) was an upper bound on the probability that $\hat{X}_i \in [x_{\min}, x_{\max}]$, the optimal value of SDP (3.39) is as well.

Appendix C

Supplemental Material for Chapter 4

C.1 Introduction

This is a condensed version of the supplementary material published to accompany the paper entitled "Dynamic Bounds on Stochastic Chemical Kinetic Systems Using Semidefinite Programming" by Garrett R. Dowdy and Paul I. Barton.

C.2 Complex Eigenvalues

In Section 4.6, we stated that it was possible to derive augmented forms of the bounding SDPs appearing in Chapter 4 that account for the imaginary components of eigenvalues. In this section, we explain how those augmented SDPs are constructed.

C.2.1 Linear Equations

Previously, we derived the equation

$$\boldsymbol{\mu}_{L}(T) - e^{\rho T} \boldsymbol{\mu}_{L}(0) = (\mathbf{A}_{L} - \rho \mathbf{I}) \int_{0}^{T} e^{\rho (T-t)} \boldsymbol{\mu}_{L}(t) dt + \mathbf{A}_{H} \int_{0}^{T} e^{\rho (T-t)} \boldsymbol{\mu}_{H}(t) dt. \quad (C.1)$$

While we derived this equation assuming that $\rho \in \mathbb{R}$, it is equally valid for any $\rho \in \mathbb{C}$, where \mathbb{C} is the set of complex numbers. Substituting $\rho = a + bi$ into the equation, we obtain

$$\begin{aligned} \boldsymbol{\mu}_{L}(T) - e^{(a+bi)T} \boldsymbol{\mu}_{L}(0) \\ &= (\mathbf{A}_{L} - (a+bi)\mathbf{I}) \int_{0}^{T} e^{(a+bi)(T-t)} \boldsymbol{\mu}_{L}(t) dt + \mathbf{A}_{H} \int_{0}^{T} e^{(a+bi)(T-t)} \boldsymbol{\mu}_{H}(t) dt \\ \boldsymbol{\mu}_{L}(T) - e^{aT} e^{biT} \boldsymbol{\mu}_{L}(0) \\ &= (\mathbf{A}_{L} - (a+bi)\mathbf{I}) \int_{0}^{T} e^{a(T-t)} e^{bi(T-t)} \boldsymbol{\mu}_{L}(t) dt + \mathbf{A}_{H} \int_{0}^{T} e^{a(T-t)} e^{bi(T-t)} \boldsymbol{\mu}_{H}(t) dt. \end{aligned}$$
(C.2)

Using Euler's formula, this becomes

$$\boldsymbol{\mu}_{L}(T) - e^{aT} (\cos(bT) + i\sin(bT)) \boldsymbol{\mu}_{L}(0)$$

$$= (\mathbf{A}_{L} - (a + bi)\mathbf{I}) \int_{0}^{T} e^{a(T-t)} (\cos(b(T-t)) + i\sin(b(T-t))) \boldsymbol{\mu}_{L}(t) dt$$

$$+ \mathbf{A}_{H} \int_{0}^{T} e^{a(T-t)} (\cos(b(T-t)) + i\sin(b(T-t))) \boldsymbol{\mu}_{H}(t) dt$$
(C.3)

Next, through simple but tedious algebraic manipulations, each side of this equation can then be separated into the real and imaginary parts:

$$(\boldsymbol{\mu}_{L}(T) - e^{aT}\cos(bT)\boldsymbol{\mu}_{L}(0)) - ie^{aT}\sin(bT)\boldsymbol{\mu}_{L}(0)$$

$$= (\mathbf{A}_{L} - a\mathbf{I})\int_{0}^{T} e^{a(T-t)}\cos(b(T-t))\boldsymbol{\mu}_{L}(t)dt$$

$$+ b\int_{0}^{T} e^{a(T-t)}\sin(b(T-t))\boldsymbol{\mu}_{L}(t)dt$$

$$+ \mathbf{A}_{H}\int_{0}^{T} e^{a(T-t)}\cos(b(T-t))\boldsymbol{\mu}_{H}(t)dt \qquad (C.4)$$

$$+ i(\mathbf{A}_{L} - a\mathbf{I})\int_{0}^{T} e^{a(T-t)}\sin(b(T-t))\boldsymbol{\mu}_{L}(t)dt$$

$$- ib\int_{0}^{T} e^{a(T-t)}\cos(b(T-t))\boldsymbol{\mu}_{L}(t)dt$$

$$+ i\mathbf{A}_{H}\int_{0}^{T} e^{a(T-t)}\sin(b(T-t))\boldsymbol{\mu}_{H}(t)dt$$

This equation holds if and only if the real part of the left-hand side equals the real part of the right-hand side, and the imaginary part of the left-hand side equals the imaginary part of the right-hand side. So, it is equivalent to the following two equations:

$$\boldsymbol{\mu}_{L}(T) - e^{aT}\cos(bT)\boldsymbol{\mu}_{L}(0) = (\mathbf{A}_{L} - a\mathbf{I})\int_{0}^{T} e^{a(T-t)}\cos(b(T-t))\boldsymbol{\mu}_{L}(t)dt$$
$$+ b\int_{0}^{T} e^{a(T-t)}\sin(b(T-t))\boldsymbol{\mu}_{L}(t)dt \qquad (C.5)$$
$$+ \mathbf{A}_{H}\int_{0}^{T} e^{a(T-t)}\cos(b(T-t))\boldsymbol{\mu}_{H}(t)dt$$

$$-e^{aT}\sin(bT)\boldsymbol{\mu}_{L}(0) = (\mathbf{A}_{L} - a\mathbf{I})\int_{0}^{T} e^{a(T-t)}\sin(b(T-t))\boldsymbol{\mu}_{L}(t)dt$$
$$-b\int_{0}^{T} e^{a(T-t)}\cos(b(T-t))\boldsymbol{\mu}_{L}(t)dt$$
$$+\mathbf{A}_{H}\int_{0}^{T} e^{a(T-t)}\sin(b(T-t))\boldsymbol{\mu}_{H}(t)dt$$
(C.6)

Now, just as we defined the variables $\mathbf{z}_{L}^{(\rho)}$ and $\mathbf{z}_{H}^{(\rho)}$ as a shorthand representation of the integrals appearing in (C.1), we will similarly define shorthand names for the integrals appearing in (C.5) and (C.6). In particular, defining

$$z_{\mathbf{j}}^{(a\cos b)} \equiv \int_{0}^{T} e^{a(T-t)} \cos(b(T-t)) \mu_{\mathbf{j}}(t) dt, \quad \forall \mathbf{j} \in \mathbb{N}^{\hat{N}},$$

$$z_{\mathbf{j}}^{(a\sin b)} \equiv \int_{0}^{T} e^{a(T-t)} \sin(b(T-t)) \mu_{\mathbf{j}}(t) dt, \quad \forall \mathbf{j} \in \mathbb{N}^{\hat{N}},$$
(C.7)

we can write the Equations (C.5) and (C.6) more concisely as

$$\boldsymbol{\mu}_{L}(T) - e^{aT}\cos(bT)\boldsymbol{\mu}_{L}(0) = (\mathbf{A}_{L} - a\mathbf{I})\mathbf{z}_{L}^{(a\cos b)} + b\mathbf{z}_{L}^{(a\sin b)} + \mathbf{A}_{H}\mathbf{z}_{H}^{(a\cos b)}$$
(C.8)

$$-e^{aT}\sin(bT)\boldsymbol{\mu}_L(0) = (\mathbf{A}_L - a\mathbf{I})\mathbf{z}_L^{(a\sin b)} - b\mathbf{z}_L^{(a\cos b)} + \mathbf{A}_H \mathbf{z}_H^{(a\sin b)}$$
(C.9)

Now for a sanity check. Since the above equations equations were derived assuming

any values $a, b \in \mathbb{R}$, they should also hold in the special case when b = 0. In this case, $\rho = a + bi = a$ is just a real number, and we would expect Equations (C.8) and (C.9) to reduce to Equation (C.1), which we derived assuming that ρ was real. Indeed, this is true. Using the definitions given in Equation (C.7), we see that Equations (C.8) reduces to Equation (C.1) and (C.9) reduces to the vacuously true statement 0 = 0. This result shows that Equations (C.8) and (C.9) are a consistent generalization of our previous work.

Of course, to use these necessary conditions as constraints in an optimization problem, we need to state them in terms of the appropriate decision variable proxies:

$$\tilde{\boldsymbol{\mu}}_{L}(T) - e^{aT}\cos(bT)\boldsymbol{\mu}_{L}(0) = (\mathbf{A}_{L} - a\mathbf{I})\tilde{\mathbf{z}}_{L}^{(a\cos b)} + b\tilde{\mathbf{z}}_{L}^{(a\sin b)} + \mathbf{A}_{H}\tilde{\mathbf{z}}_{H}^{(a\cos b)}$$
(C.10)

$$-e^{aT}\sin(bT)\boldsymbol{\mu}_L(0) = (\mathbf{A}_L - a\mathbf{I})\tilde{\mathbf{z}}_L^{(a\sin b)} - b\tilde{\mathbf{z}}_L^{(a\cos b)} + \mathbf{A}_H\tilde{\mathbf{z}}_H^{(a\sin b)}$$
(C.11)

C.2.2 Additional Constraints

As was the case with Equation (C.1), Equations (C.10) and (C.11) are of limited value unless we further constrain the values of $\tilde{z}_{j}^{(a\cos b)}$ and $\tilde{z}_{j}^{(a\cos b)}$.

Zero-Order Moments

The first set of these additional constraints is derived from the fact that $\mu_0(t) = 1$ throughout time. Using this fact with Equation (C.7) gives

$$z_{0}^{(a\cos b)} = \int_{0}^{T} e^{a(T-t)} \cos(b(T-t))\mu_{0}(t)dt,$$

=
$$\int_{0}^{T} e^{a(T-t)} \cos(b(T-t))dt,$$

=
$$\frac{e^{aT}(b\sin(bT) + a\cos(bT)) - a}{a^{2} + b^{2}},$$
 (C.12)

and

$$z_{0}^{(a\sin b)} = \int_{0}^{T} e^{a(T-t)} \sin(b(T-t))\mu_{0}(t)dt$$

=
$$\int_{0}^{T} e^{a(T-t)} \sin(b(T-t))dt,$$

=
$$\frac{e^{aT}(a\sin(bT) - b\cos(bT)) + b}{a^{2} + b^{2}}.$$
 (C.13)

Writing these equations for the proxy variables gives:

$$\tilde{z}_{0}^{(a\cos b)} = \frac{e^{aT}(b\sin(bT) + a\cos(bT)) - a}{a^2 + b^2},$$
(C.14)

and

$$\tilde{z}_{0}^{(a\sin b)} = \frac{e^{aT}(a\sin(bT) - b\cos(bT)) + b}{a^{2} + b^{2}}.$$
(C.15)

LMI and Second-Order Cone Constraints

The second set of additional constraints will take the form of LMIs and second-order cone constraints.

First, by the trigonometric identity, we have

$$\sin^2(b(T-t)) + \cos^2(b(T-t)) = 1, \quad \forall t \in [0,T].$$
(C.16)

It follows trivially, that

$$\sin^2(b(T-t)) + \cos^2(b(T-t)) \le 1, \quad \forall t \in [0,T].$$
(C.17)

Now, pick an arbitrary $\mathbf{j} \in \mathbb{N}^{\hat{N}}$. Multiplying both sides of the above inequality by $\mu_{\mathbf{j}}^2(t)$ gives

$$\mu_{\mathbf{j}}^{2}(t)\sin^{2}(b(T-t)) + \mu_{\mathbf{j}}^{2}(t)\cos^{2}(b(T-t)) \le \mu_{\mathbf{j}}^{2}(t), \quad \forall t \in [0,T].$$
(C.18)

Taking the square root of both sides gives

$$\sqrt{\mu_{\mathbf{j}}^{2}(t)\sin^{2}(b(T-t)) + \mu_{\mathbf{j}}^{2}(t)\cos^{2}(b(T-t))} = \left\| \begin{bmatrix} \mu_{\mathbf{j}}(t)\sin(b(T-t)) \\ \mu_{\mathbf{j}}(t)\cos(b(T-t)) \end{bmatrix} \right\|_{2}, \quad (C.19)$$

$$\leq |\mu_{\mathbf{j}}(t)|, \\
= \mu_{\mathbf{j}}(t).$$

for all $t \in [0, T]$. The last equality assumes that $\mu_{\mathbf{j}}(t) \geq 0$. This will be true as long as we choose the appropriate representation of our state space. It is true, in particular, for our chosen representation in terms of independent species.

Multiplying both sides of the above inequality by $e^{a(T-t)}$ and integrating from t = 0 to t = T gives

$$\int_0^T e^{a(T-t)} \left\| \left[\begin{array}{c} \mu_{\mathbf{j}}(t)\sin(b(T-t))\\ \mu_{\mathbf{j}}(t)\cos(b(T-t)) \end{array} \right] \right\|_2 dt \le \int_0^T e^{a(T-t)}\mu_{\mathbf{j}}(t)dt = z_{\mathbf{j}}^{(a)} \tag{C.20}$$

Focusing on the left-hand side, we can use Jensen's Inequality to obtain

$$\begin{split} \int_{0}^{T} e^{a(T-t)} \left\| \begin{bmatrix} \mu_{\mathbf{j}}(t) \sin(b(T-t)) \\ \mu_{\mathbf{j}}(t) \cos(b(T-t)) \end{bmatrix} \right\|_{2} dt &= \int_{0}^{T} \left\| \begin{bmatrix} e^{a(T-t)}\mu_{\mathbf{j}}(t) \sin(b(T-t)) \\ e^{a(T-t)}\mu_{\mathbf{j}}(t) \cos(b(T-t)) \end{bmatrix} \right\|_{2} dt \\ &\geq \left\| \int_{0}^{T} \begin{bmatrix} e^{a(T-t)}\mu_{\mathbf{j}}(t) \sin(b(T-t)) \\ e^{a(T-t)}\mu_{\mathbf{j}}(t) \cos(b(T-t)) \end{bmatrix} dt \right\|_{2} \\ &= \left\| \begin{bmatrix} \int_{0}^{T} e^{a(T-t)}\mu_{\mathbf{j}}(t) \sin(b(T-t)) dt \\ \int_{0}^{T} e^{a(T-t)}\mu_{\mathbf{j}}(t) \cos(b(T-t)) dt \end{bmatrix} \right\|_{2} \\ &= \left\| \begin{bmatrix} z_{\mathbf{j}}^{(a \sin b)} \\ z_{\mathbf{j}}^{(a \cos b)} \end{bmatrix} \right\|_{2} \end{split}$$
(C.21)

Combining this result with Inequality (C.20), we have the following second-order cone

constraint involving the quantities $z_{\mathbf{j}}^{(a \sin b)}$, $z_{\mathbf{j}}^{(a \cos b)}$, and $z_{\mathbf{j}}^{(a)}$.

$$\left\| \begin{bmatrix} z_{\mathbf{j}}^{(a\sin b)} \\ z_{\mathbf{j}}^{(a\cos b)} \end{bmatrix} \right\|_{2} \le z_{\mathbf{j}}^{(a)}, \tag{C.22}$$

which, of course, can also be written for our proxy variables $\tilde{z}_{\mathbf{j}}^{(a \operatorname{cos} b)}$, $\tilde{z}_{\mathbf{j}}^{(a \cos b)}$, and $\tilde{z}_{\mathbf{j}}^{(a)}$.

Now, unless we constrain the value of $\tilde{z}_{\mathbf{j}}^{(a)}$, this inequality is of no use to us; if $\tilde{z}_{\mathbf{j}}^{(a)}$ can be made arbitrarily large, then $\tilde{z}_{\mathbf{j}}^{(a \sin b)}$ and $\tilde{z}_{\mathbf{j}}^{(a \cos b)}$ can take any value. Fortunately, we already know exactly how to constrain $\tilde{z}_{\mathbf{j}}^{(a)}$. It is subject to exactly the same constraints that we derived earlier for $\tilde{z}_{\mathbf{j}}^{(\rho)}$, when we were considering exclusively real values of ρ . First, we have the linear equalities,

$$\tilde{\boldsymbol{\mu}}_L(T) - e^{aT} \boldsymbol{\mu}_L(0) = (\mathbf{A}_L - a\mathbf{I})\tilde{\mathbf{z}}_L^{(a)} + \mathbf{A}_H \tilde{\mathbf{z}}_H^{(a)}.$$
 (C.23)

Second, we have the LMIs implied by membership in the cone $C_n(\boldsymbol{\alpha}, \boldsymbol{\beta})$,

$$\tilde{\mathbf{z}}^{(a)} \in C_n(\boldsymbol{\alpha}, \boldsymbol{\beta})$$
 (C.24)

C.2.3 Bringing It All Together

To summarize, if you wish to use a complex $\rho = a + bi$ in the bound calculation, we need to include the following constraints in the optimization problem:

$$\tilde{\boldsymbol{\mu}}_{L}(T) - e^{aT}\cos(bT)\boldsymbol{\mu}_{L}(0) = (\mathbf{A}_{L} - a\mathbf{I})\tilde{\mathbf{z}}_{L}^{(a\cos b)} + b\tilde{\mathbf{z}}_{L}^{(a\sin b)} + \mathbf{A}_{H}\tilde{\mathbf{z}}_{H}^{(a\cos b)}$$
(C.25)

$$-e^{aT}\sin(bT)\boldsymbol{\mu}_L(0) = (\mathbf{A}_L - a\mathbf{I})\tilde{\mathbf{z}}_L^{(a\sin b)} - b\tilde{\mathbf{z}}_L^{(a\cos b)} + \mathbf{A}_H\tilde{\mathbf{z}}_H^{(a\sin b)}$$
(C.26)

$$\tilde{z}_{0}^{(a\cos b)} = \frac{e^{aT}(b\sin(bT) + a\cos(bT)) - a}{a^2 + b^2},$$
(C.27)

$$\tilde{z}_{\mathbf{0}}^{(a\sin b)} = \frac{e^{aT}(a\sin(bT) - b\cos(bT)) + b}{a^2 + b^2}.$$
(C.28)

$$\left\| \begin{bmatrix} \tilde{z}_{\mathbf{j}}^{(a \sin b)} \\ \tilde{z}_{\mathbf{j}}^{(a \cos b)} \end{bmatrix} \right\|_{2} \leq \tilde{z}_{\mathbf{j}}^{(a)}, \quad \forall \mathbf{j} \in \mathbb{N}^{\hat{N}} \text{s.t. } |\mathbf{j}| \leq M$$
(C.29)

$$\tilde{\boldsymbol{\mu}}_{L}(T) - e^{aT} \boldsymbol{\mu}_{L}(0) = (\mathbf{A}_{L} - \rho \mathbf{I}) \tilde{\mathbf{z}}_{L}^{(a)} + \mathbf{A}_{H} \tilde{\mathbf{z}}_{H}^{(a)}$$
(C.30)

$$\tilde{\mathbf{z}}^{(a)} \in C_n(\boldsymbol{\alpha}, \boldsymbol{\beta})$$
 (C.31)

As a reminder, the M appearing in Constraint (C.29) is the order of the highest-order moment appearing in $\tilde{\mathbf{z}}_{H}^{(a)}$ (and $\boldsymbol{\mu}_{H}$).

C.2.4 Complex Conjugates

Now, recall that the values of ρ that we use are intended to estimate eigenvalues of the infinitesimal generator matrix **G**. All of the elements of this matrix are real. This implies that any complex eigenvalues occur in conjugate pairs. Thus, if we are using $\rho = a + bi$ as an estimate of one of the eigenvalues, it seems that we should be using $\rho = a - bi$ also. However, one can show that if we write out Conditions (C.25) -(C.31) for $\rho = a + bi$, it is redundant to also write out these conditions for $\rho = a - bi$.

To see this, suppose we write out

$$\tilde{\boldsymbol{\mu}}_{L}(T) - e^{aT}\cos(bT)\boldsymbol{\mu}_{L}(0) = (\mathbf{A}_{L} - a\mathbf{I})\tilde{\mathbf{z}}_{L}^{(a\cos-b)} - b\tilde{\mathbf{z}}_{L}^{(a\sin-b)} + \mathbf{A}_{H}\tilde{\mathbf{z}}_{H}^{(a\cos-b)}, \quad (C.32)$$

$$e^{aT}\sin(bT)\boldsymbol{\mu}_L(0) = (\mathbf{A}_L - a\mathbf{I})\tilde{\mathbf{z}}_L^{(a\sin-b)} + b\tilde{\mathbf{z}}_L^{(a\cos-b)} + \mathbf{A}_H\tilde{\mathbf{z}}_H^{(a\sin-b)}, \qquad (C.33)$$

$$\tilde{z}_{0}^{(a\cos-b)} = \frac{e^{aT}(b\sin(bT) + a\cos(bT)) - a}{a^2 + b^2},$$
(C.34)

$$\tilde{z}_{0}^{(a\sin-b)} = \frac{e^{aT}(-a\sin(bT) + b\cos(bT)) - b}{a^2 + b^2}.$$
(C.35)

$$\left\| \begin{bmatrix} \tilde{z}_{\mathbf{j}}^{(a\sin-b)} \\ \tilde{z}_{\mathbf{j}}^{(a\cos-b)} \end{bmatrix} \right\|_{2} \le \tilde{z}_{\mathbf{j}}^{(a)}, \quad \forall \mathbf{j} \in \mathbb{N}^{\hat{N}} \text{s.t. } |\mathbf{j}| \le M.$$
(C.36)

to supplement Conditions (C.25) - (C.31) as constraints in our optimization problem. (Conditions (C.30)- (C.31) are not copied, as they have no dependence on the sign of b, and are thus obviously redundant.) The question is: does the addition of Constraints (C.32) - (C.36) further constrain the set of feasible vectors $\tilde{\mu}_L(T)$? The answer is "no", because if we have vectors $\tilde{\mu}_L(T)$, $\tilde{\mathbf{z}}^{(a\cos b)}$, $\tilde{\mathbf{z}}^{(a\sin b)}$, and $\tilde{\mathbf{z}}^{(a)}$ which satisfy Constraints (C.25) - (C.31), we can trivially construct vectors $\tilde{\mathbf{z}}^{(a\cos -b)}$ and $\tilde{\mathbf{z}}^{(a\sin -b)}$ which satisfy Constraints (C.32) - (C.36). We simply set $\tilde{\mathbf{z}}^{(a\cos -b)} \equiv \tilde{\mathbf{z}}^{(a\cos b)}$ and $\tilde{\mathbf{z}}^{(a\sin -b)} \equiv -\tilde{\mathbf{z}}^{(a\sin b)}$. This means that the addition of Constraints (C.32) - (C.36) has no effect on the set of feasible vectors $\tilde{\boldsymbol{\mu}}_L(T)$. In other words, Constraints (C.32) - (C.36) are redundant.

From another perspective, this conclusion is not surprising. If we start with Conditions (C.25) - (C.31), written for the true quantities $\boldsymbol{\mu}_L(T)$, $\mathbf{z}^{(a\cos b)}$, $\mathbf{z}^{(a\sin b)}$, and $\mathbf{z}^{(a)}$, and we expand the definitions of $\mathbf{z}^{(a\cos b)}$ and $\mathbf{z}^{(a\sin b)}$ into integral form, we see that replacing b with -b yields an exactly equivalent set of conditions.

To summarize, if we write Conditions (C.25) - (C.31) for $\rho = a + bi$, we don't need to worry about also writing them for $\rho = a - bi$ as the second set of conditions is implied by the first.

C.2.5 An Augmented SDP

So suppose we wanted to construct an SDP for bounding a stochastic chemical kinetic system using a set \mathcal{R} containing both real and complex values of ρ . What does this look like?

Suppose that there are $|\mathcal{R}|$ values of ρ , which can be written as $\mathcal{R} \equiv \{\rho_1, \ldots, \rho_{|\mathcal{R}|}\} \equiv \{a_1 + b_1 i, \ldots, a_{|\mathcal{R}|} + b_{|\mathcal{R}|} i\}$. For the reasons discussed in the previous section we do

not want this set to contain any complex conjugate pairs. Let $J \equiv \{1, \ldots, |\mathcal{R}|\}$ and let $J_{\mathbb{C}} \equiv \{j \in J : b_j \neq 0\}$. In other words, $J_{\mathbb{C}}$ is the set of indices corresponding to the complex values of ρ_j .

Then, we can write the SDP for calculating an upper bound on the mean molecular count of species i as follows:

$$\begin{split} \max_{\substack{\boldsymbol{\mu}(T),\\\boldsymbol{\mu}(T),\\\boldsymbol{\mu}(T)\in J,\\\boldsymbol{\mu}(T)\in J_{C},\\\boldsymbol{\mu}(a)\cosh^{j},\forall j\in J_{C},\\\boldsymbol{\mu}(T)\in C_{n}(\boldsymbol{\alpha},\boldsymbol{\beta}),\\ \boldsymbol{\bar{x}}^{(a_{j})}\in C_{n}(\boldsymbol{\alpha},\boldsymbol{\beta}), \quad \forall j\in J,\\ \boldsymbol{\bar{\mu}}(T)\in C_{n}(\boldsymbol{\alpha},\boldsymbol{\beta}), \quad \forall j\in J,\\ \boldsymbol{\bar{\mu}}(T)=C_{n}(\boldsymbol{\alpha},\boldsymbol{\beta}), \quad \forall j\in J,\\ \boldsymbol{\bar{\mu}}_{L}(T)=e^{a_{j}T}\boldsymbol{\mu}_{L}(0)=(\mathbf{A}_{L}-a_{j}\mathbf{I})\boldsymbol{\bar{x}}_{L}^{(a_{j})}+\mathbf{A}_{H}\boldsymbol{\bar{x}}_{H}^{(a_{j})}, \quad \forall j\in J,\\ \boldsymbol{\bar{x}}_{0}^{(a_{j})}=\begin{cases} T & \text{if } a_{j}=0,\\ \frac{e^{a_{j}T}-1}{a_{j}} & \text{otherwise}, \end{cases} \quad \forall j\in J,\\ \boldsymbol{\bar{\mu}}_{L}(T)=e^{a_{j}T}\cos(b_{j}T)\boldsymbol{\mu}_{L}(0)=\\ & (\mathbf{A}_{L}-a_{j}\mathbf{I})\boldsymbol{\bar{x}}_{L}^{(a_{j}\cos b_{j})}+b_{j}\boldsymbol{\bar{x}}_{L}^{(a_{j}\sin b_{j})}+\mathbf{A}_{H}\boldsymbol{\bar{x}}_{H}^{(a_{j}\cos b_{j})}, \quad \forall j\in J_{C},\\ & -e^{a_{j}T}\sin(b_{j}T)\boldsymbol{\mu}_{L}(0)=\\ & (\mathbf{A}_{L}-a_{j}\mathbf{I})\boldsymbol{\bar{x}}_{L}^{(a_{j}\cos b_{j})}-b_{j}\boldsymbol{\bar{x}}_{L}^{(a_{j}\cos b_{j})}+\mathbf{A}_{H}\boldsymbol{\bar{x}}_{H}^{(a_{j}\sin b_{j})}, \quad \forall j\in J_{C},\\ & \boldsymbol{\bar{z}}_{0}^{(a\cos b)}&=\frac{e^{aT}(b\sin(bT)+a\cos(bT))-a}{a^{2}+b^{2}}, \quad \forall j\in J_{C},\\ & \boldsymbol{\bar{z}}_{0}^{(a\cos b)}&=\frac{e^{aT}(a\sin(bT)-b\cos(bT))+b}{a^{2}+b^{2}}, \quad \forall j\in J_{C},\\ & \left\|\left\|\begin{bmatrix}\boldsymbol{\bar{x}}_{1}^{(a_{j}\sin b_{j})}\\\boldsymbol{\bar{x}_{j}^{(a_{j}\cos b_{j})}\end{bmatrix}\right\|\right\|_{2}\leq \boldsymbol{\bar{x}}_{0}^{(a_{j})}, \quad \forall \mathbf{j}\in\mathbb{N}^{\hat{N}}\text{s.t.} |\mathbf{j}|\leq M, \quad \forall j\in J_{C}.\\ & (C.37) \end{split}$$

Technically, this is not an SDP, because of the presence of the second order cone constraints. However, each of these second order cone constraints can be written equivalently as an LMI, so this technicality is of no consequence.

C.3 The Absence of the Closure Problem

In Section 4.7, we observed that when our bounding method is applied to systems which do not exhibit the closure problem, the bounds are often perfect. In this section, we give some insight into why this is the case.

C.3.1 Theoretical Reasoning

When a system does not suffer from the closure problem, the \mathbf{A}_H matrix appearing in Equation (C.1) is all zeros. This means that the associated constraint simplifies to

$$\tilde{\boldsymbol{\mu}}_L(T) - e^{\rho T} \boldsymbol{\mu}_L(0) = (\mathbf{A}_L - \rho \mathbf{I}) \tilde{\mathbf{z}}_L^{(\rho)}.$$
(C.38)

Now, to simplify the discussion, let us assume that the matrix has $\mathbf{A}_L \in \mathbb{R}^{d \times d}$ has d distinct eigenvalues $\{\lambda_1, \ldots, \lambda_d\}$. It follows that \mathbf{A}_L has d linearly independent eigenvectors $\{\mathbf{v}_1, \ldots, \mathbf{v}_d\} \subset \mathbb{C}^d$. Furthermore, if we let $\mathbf{V} \equiv [\mathbf{v}_1 \ldots \mathbf{v}_d]$, then

$$\mathbf{A}_L = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^{-1}, \tag{C.39}$$

where $\Lambda \equiv \operatorname{diag}(\lambda_1, \ldots, \lambda_d)$. It follows that

$$\mathbf{A}_{L} - \rho \mathbf{I} = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^{-1} - \rho \mathbf{V} \mathbf{V}^{-1} = \mathbf{V} (\mathbf{\Lambda} - \rho \mathbf{I}) \mathbf{V}^{-1}.$$
 (C.40)

Then, Equation (C.38) becomes

$$\tilde{\boldsymbol{\mu}}_L(T) - e^{\rho T} \boldsymbol{\mu}_L(0) = \mathbf{V}(\boldsymbol{\Lambda} - \rho \mathbf{I}) \mathbf{V}^{-1} \tilde{\mathbf{z}}_L^{(\rho)}.$$
 (C.41)

Rearranging gives

$$\tilde{\boldsymbol{\mu}}_L(T) = \mathbf{V}(\boldsymbol{\Lambda} - \rho \mathbf{I})\mathbf{V}^{-1}\tilde{\mathbf{z}}_L^{(\rho)} + e^{\rho T}\boldsymbol{\mu}_L(0).$$
(C.42)

Now, suppose that we choose our values of ρ to be $\mathcal{R} = \{\lambda_1, \ldots, \lambda_d\}$. Then, enforcing Equation (C.42) for each $\rho \in \mathcal{R}$, we have

$$\tilde{\boldsymbol{\mu}}_{L}(T) = \mathbf{V}(\boldsymbol{\Lambda} - \lambda_{1}\mathbf{I})\mathbf{V}^{-1}\tilde{\mathbf{z}}_{L}^{(\lambda_{1})} + e^{\lambda_{1}T}\boldsymbol{\mu}_{L}(0),$$

$$\tilde{\boldsymbol{\mu}}_{L}(T) = \mathbf{V}(\boldsymbol{\Lambda} - \lambda_{2}\mathbf{I})\mathbf{V}^{-1}\tilde{\mathbf{z}}_{L}^{(\lambda_{2})} + e^{\lambda_{2}T}\boldsymbol{\mu}_{L}(0),$$

$$\vdots$$

$$\tilde{\boldsymbol{\mu}}_{L}(T) = \mathbf{V}(\boldsymbol{\Lambda} - \lambda_{d}\mathbf{I})\mathbf{V}^{-1}\tilde{\mathbf{z}}_{L}^{(\lambda_{d})} + e^{\lambda_{d}T}\boldsymbol{\mu}_{L}(0),$$
(C.43)

The first set of these equations specifies that $\tilde{\mu}_L(T)$ is contained in an affine subspace spanned by the vectors $\{\mathbf{v}_2, \ldots, \mathbf{v}_d\}$, with \mathbf{v}_1 omitted:

$$A_1 \equiv \left\{ \mathbf{V}(\mathbf{\Lambda} - \lambda_1 \mathbf{I}) \mathbf{V}^{-1} \mathbf{z} + e^{\lambda_1 T} \boldsymbol{\mu}_L(0) : \mathbf{z} \in \mathbb{R}^d \right\}.$$
 (C.44)

Similarly, the second set of equations specifies that $\tilde{\mu}_L(T)$ is contained in an affine subspace spanned by the vectors $\{\mathbf{v}_1, \mathbf{v}_3, \dots, \mathbf{v}_d\}$, with \mathbf{v}_2 omitted:

$$A_2 \equiv \left\{ \mathbf{V}(\mathbf{\Lambda} - \lambda_2 \mathbf{I}) \mathbf{V}^{-1} \mathbf{z} + e^{\lambda_2 T} \boldsymbol{\mu}_L(0) : \mathbf{z} \in \mathbb{R}^d \right\}$$
(C.45)

Continuing in this way, we see that $\tilde{\mu}_L(T)$ is contained in all affine subspaces

$$A_j \equiv \left\{ \mathbf{V}(\mathbf{\Lambda} - \lambda_j \mathbf{I}) \mathbf{V}^{-1} \mathbf{z} + e^{\lambda_j T} \boldsymbol{\mu}_L(0) : \mathbf{z} \in \mathbb{R}^d \right\}$$
(C.46)

for $j \in \{1, \ldots, d\}$. This implies that $\tilde{\boldsymbol{\mu}}_L(T)$ must lie in the intersection $\bigcap_{j \in \{1, \ldots, d\}} A_j$. One can show that this intersection is nonempty and has exactly one point.

Claim 9. The set $\bigcap_{j \in \{1,...,d\}} A_j$ is nonempty.

Proof. The statement that $\bigcap_{j \in \{1,...,d\}} A_j$ is nonempty is equivalent to the statement that there exists some $(\tilde{\boldsymbol{\mu}}_L(T), \tilde{\mathbf{z}}_L^{(\lambda_1)}, \ldots, \tilde{\mathbf{z}}_L^{(\lambda_d)}) \in \mathbb{R}^{d+d^2}$ satisfying Equation (C.43).

This equation can be written equivalently in matrix-vector form:

$$\begin{bmatrix} \mathbf{I} & -\mathbf{V}(\mathbf{\Lambda} - \lambda_{1}\mathbf{I})\mathbf{V}^{-1} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{I} & \mathbf{0} & -\mathbf{V}(\mathbf{\Lambda} - \lambda_{2}\mathbf{I})\mathbf{V}^{-1} & \dots & \mathbf{0} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{I} & \mathbf{0} & \mathbf{0} & \dots & -\mathbf{V}(\mathbf{\Lambda} - \lambda_{d}\mathbf{I})\mathbf{V}^{-1} \end{bmatrix} \begin{bmatrix} \tilde{\boldsymbol{\mu}}_{L}(T) \\ \tilde{\mathbf{z}}_{L}^{(\lambda_{1})} \\ \vdots \\ \tilde{\mathbf{z}}_{L}^{(\lambda_{2})} \\ \vdots \\ \tilde{\mathbf{z}}_{L}^{(\lambda_{d})} \end{bmatrix} \\ = \begin{bmatrix} e^{\lambda_{2}T}\boldsymbol{\mu}_{L}(0) \\ e^{\lambda_{2}T}\boldsymbol{\mu}_{L}(0) \\ \vdots \\ e^{\lambda_{d}T}\boldsymbol{\mu}_{L}(0) \end{bmatrix}$$
(C.47)

This system of equations is guaranteed to have a solution if the leftmost matrix (call it $\mathbf{Q} \in \mathbb{R}^{d^2 \times (d^2+d)}$) has d^2 linearly independent columns (i.e., if it's rank is d^2). To help us analyze its rank, we will pre-multiply by the invertible matrix

$$\begin{bmatrix} \mathbf{V}^{-1} & & \\ & \ddots & \\ & & \mathbf{V}^{-1} \end{bmatrix} \in \mathbb{R}^{d^2},$$
(C.48)

and post-multiply by the invertible matrix

$$\begin{bmatrix} \mathbf{V} & & \\ & \ddots & \\ & & \mathbf{V} \end{bmatrix} \in \mathbb{R}^{d^2 + d}.$$
(C.49)

The resulting matrix, which has the same rank as $\mathbf{Q},$ is

$$\begin{bmatrix} \mathbf{I} & -(\mathbf{\Lambda} - \lambda_1 \mathbf{I}) & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{I} & \mathbf{0} & -(\mathbf{\Lambda} - \lambda_2 \mathbf{I}) & \dots & \mathbf{0} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{I} & \mathbf{0} & \mathbf{0} & \dots & -(\mathbf{\Lambda} - \lambda_d \mathbf{I}) \end{bmatrix}.$$
 (C.50)

From the sparsity pattern of this matrix, it follows that no row can be expressed as a linear combination of the others. It follows that each of the d^2 rows are linearly independent. This implies that the rank of **Q** is d^2 , and we are done.

Claim 10. The set $\bigcap_{j \in \{1,...,d\}} A_j$ consists of a unique point.

Proof. First, some preliminaries. Let $\{\mathbf{r}_1^{\mathrm{T}}, \ldots, \mathbf{r}_d^{\mathrm{T}}\}$ denote the rows of the matrix \mathbf{V}^{-1} , so that

$$\begin{bmatrix} \mathbf{r}_1^{\mathrm{T}} \\ \vdots \\ \mathbf{r}_d^{\mathrm{T}} \end{bmatrix} \equiv \mathbf{V}^{-1}$$
(C.51)

Since \mathbf{V}^{-1} is invertible, it follows that the vectors $\{\mathbf{r}_1, \ldots, \mathbf{r}_d\} \subset \mathbb{R}^d$ are linearly independent. Furthermore, since $\mathbf{V}^{-1}\mathbf{V} = \mathbf{I}$, it follows that $\mathbf{r}_j^{\mathrm{T}}\mathbf{V} = \mathbf{e}_j^{\mathrm{T}}$ for all $j \in$ $\{1, \ldots, d\}$. Finally, for all $j \in \{1, \ldots, d\}$, and for all $\mathbf{x}^{(1)}, \mathbf{x}^{(2)} \in A_j$, it follows that the difference $\mathbf{x}^{(1)} - \mathbf{x}^{(2)}$ is orthogonal to \mathbf{r}_j . To see this consider

$$\mathbf{r}_{j}^{\mathrm{T}}(\mathbf{x}^{(1)} - \mathbf{x}^{(2)})$$

$$= \mathbf{r}_{j}^{\mathrm{T}}\left(\left(\mathbf{V}(\mathbf{\Lambda} - \lambda_{j}\mathbf{I})\mathbf{V}^{-1}\mathbf{z}^{(1)} + e^{\lambda_{j}T}\boldsymbol{\mu}_{L}(0)\right) - \left(\mathbf{V}(\mathbf{\Lambda} - \lambda_{j}\mathbf{I})\mathbf{V}^{-1}\mathbf{z}^{(2)} + e^{\lambda_{j}T}\boldsymbol{\mu}_{L}(0)\right)\right)$$

$$= \mathbf{r}_{j}^{\mathrm{T}}\mathbf{V}(\mathbf{\Lambda} - \lambda_{j}\mathbf{I})\mathbf{V}^{-1}(\mathbf{z}^{(1)} - \mathbf{z}^{(2)})$$

$$= \mathbf{e}_{j}^{\mathrm{T}}(\mathbf{\Lambda} - \lambda_{j}\mathbf{I})\mathbf{V}^{-1}(\mathbf{z}^{(1)} - \mathbf{z}^{(2)})$$

$$= \mathbf{0}^{\mathrm{T}}\mathbf{V}^{-1}(\mathbf{z}^{(1)} - \mathbf{z}^{(2)})$$

$$= 0$$
(C.52)

Now, we have already established that $\bigcap_{j \in \{1,...,d\}} A_j$ contains at least one point \mathbf{x} . Suppose there was some $\mathbf{x}' \in \mathbb{R}^d$ such that $\mathbf{x}' \neq \mathbf{x}$ which was also in $\bigcap_{j \in \{1,...,d\}} A_j$. From Equation (C.52), it follows that $\mathbf{x} - \mathbf{x}'$ is orthogonal to all \mathbf{r}_j , where $j \in \{1,...,d\}$. This can equivalently be written as $\mathbf{V}^{-1}(\mathbf{x} - \mathbf{x}') = \mathbf{0}$. Since \mathbf{V}^{-1} is invertible, this implies that $\mathbf{x} - \mathbf{x}' = \mathbf{0}$ or that $\mathbf{x} = \mathbf{x}'$, which is a contradiction. It follows that our of assumption of the distinct $\mathbf{x}' \in \bigcap_{j \in \{1,...,d\}} A_j$ was false and that \mathbf{x} is the unique point in $\bigcap_{j \in \{1,...,d\}} A_j$.

Now, what does all this mean? If we include Equations (C.43) as constraints in

SDP (4.41), we are restricting the set of feasible $\tilde{\boldsymbol{\mu}}_L(T)$ vectors to at most one point $\tilde{\boldsymbol{\mu}}_L(T)^*$. As long as the other constraints (e.g., the LMIs) do not exclude $\tilde{\boldsymbol{\mu}}_L(T)^*$, it will be the only feasible $\tilde{\boldsymbol{\mu}}_L(T)$ vector for the problem. Since the objective function depends only on $\tilde{\boldsymbol{\mu}}_L(T)$, it then doesn't matter whether we attempt to maximize or minimize the objective function, the optimal value will simply be the objective function function evaluated at $\tilde{\boldsymbol{\mu}}_L(T)^*$. This means that the upper and lower bounds that we calculate, in theory, will be the same.

The astute reader might find it odd that, in the analysis above, we have chosen our set \mathcal{R} to match the eigenvalues of the matrix \mathbf{A}_L , while in Section 4.3.7 we identified the smallest-magnitude eigenvalues of the infinitesimal generator matrix \mathbf{G} as being important. This apparent disconnect is resolved by the observation that, in the absence of the closure problem, there is a close connection between the eigenvalues of the matrices \mathbf{G} and those of \mathbf{A}_L . We are not able to state this connection with mathematical precision at present. However, we have noticed, for several such problems, that the smallest magnitude eigenvalues of \mathbf{G} are equal to those of \mathbf{A}_L .

C.4 Bounds on Higher-Order Moments

In this section, we provide a brief demonstration of the bounding method's ability to calculate bounds on higher order moments. This demonstration is based on the reaction system from Section 4.5, namely:

$$A + B \xrightarrow{c_1} C \xleftarrow{c_2} C_3 D$$

We will consider species A and C as the independent species, and we will bound the moments $\mu_{(0,3)}(t) = \langle x_c^3(t) \rangle$ and $\mu_{(2,2)}(t) = \langle x_A^2 x_C^2(t) \rangle$. This is done by solving a minor modification of the SDP for bounding the mean molecular count of species *i*, in which the objective function is changed to $\tilde{\mu}_{(0,3)}(T)$ and $\tilde{\mu}_{(2,2)}(T)$, in turn. Solving the maximization and minimization version of this modified SDP with m = 3 and $\mathcal{R} = \{0, -2, -2.4, -4.4\}$ gives the bounds shown in Figures C-1 and C-2 below. For comparison, the actual trajectories of these moments were computed by direct solution of the CME and plotted along with the bounds.



Figure C-1: Time-varying bounds on a third-order moment.



Figure C-2: Time-varying bounds on a fourth-order moment.

C.5 Proof of the Monotonicity of the Bounds with Increasing m

In Section 4.4.2, we asserted that "increasing the value of the parameter m leads to monotonically tighter bounds". We will now rigorously state and prove this claim.

Claim 11. Let $\langle \hat{X}_i(T) \rangle_{m_1}^U$ be the upper bound obtained on the mean molecular count of independent species *i* obtained by solving SDP (4.41) with a given set \mathcal{R} and $m = m_1 \in \mathbb{N}$. Similarly, let $\langle \hat{X}_i(T) \rangle_{m_2}^U$ be the upper bound obtained on the mean molecular count of independent species *i* obtained by solving SDP (4.41) with the same set \mathcal{R} and with $m = m_2 \in \mathbb{N}$ such that $m_2 = m_1 + 1$. Then $\langle \hat{X}_i(T) \rangle_{m_1}^U \geq \langle \hat{X}_i(T) \rangle_{m_2}^U$.

Proof. Let SDP_j denote the SDP solved to obtain $\langle \hat{X}_i(T) \rangle_{m_j}^U$. The essential idea of this proof is to show that for every feasible point of SDP_2 we can construct a feasible point of SDP_1 that achieves the same objective value. This is enough to imply the conclusion of our claim.

First, we know that SDP_2 must have at least one feasible point $(\tilde{\boldsymbol{\mu}}(T), (\tilde{\mathbf{z}}^{(\rho)})_{\rho \in \mathcal{R}})$. By the design of SDP_2 , the true moment vector $(\boldsymbol{\mu}(T), (\mathbf{z}^{(\rho)})_{\rho \in \mathcal{R}})$ associated with the true time-varying probability distribution, $P(\cdot, t), t \in [0, T]$ satisfies all the necessary conditions which act as constraints. Thus, this point is feasible for SDP_2 .

Let $(\tilde{\boldsymbol{\mu}}(T), (\tilde{\mathbf{z}}^{(\rho)})_{\rho \in \mathcal{R}})$ be an arbitrary feasible point for SDP₂. By the definition of feasibility, the point $(\tilde{\boldsymbol{\mu}}(T), (\tilde{\mathbf{z}}^{(\rho)})_{\rho \in \mathcal{R}})$ must satisfy all the constraints appearing in SDP₂. In particular, it must satisfy

$$\tilde{\mu}_0(T) = 1 \tag{C.53}$$

$$\tilde{\boldsymbol{\mu}}(T) \in C_{n_2}(\boldsymbol{\alpha}, \boldsymbol{\beta})$$
 (C.54)

$$\tilde{\mathbf{z}}^{(\rho)} \in C_{n_2}(\boldsymbol{\alpha}, \boldsymbol{\beta}), \quad \forall \rho \in \mathcal{R}$$
 (C.55)

$$\tilde{\boldsymbol{\mu}}_{L}(T) - e^{\rho T} \boldsymbol{\mu}_{L}(0) = (\mathbf{A}_{L}^{(2)} - \rho \mathbf{I}) \tilde{\mathbf{z}}_{L}^{(\rho)} + \mathbf{A}_{H}^{(2)} \tilde{\mathbf{z}}_{H}^{(\rho)}, \quad \forall \rho \in \mathcal{R}$$
(C.56)

$$\tilde{z}_{\mathbf{0}}^{(\rho)} = \begin{cases} T & \text{if } \rho = 0, \\ \frac{e^{\rho T} - 1}{\rho} & \text{otherwise} \end{cases} \quad \forall \rho \in \mathcal{R},$$
(C.57)

where the matrices $\mathbf{A}_{L}^{(2)}$ and $\mathbf{A}_{H}^{(2)}$ contain rows corresponding to all moments up to order m_2 , and where

$$n_2 = \left\lceil \frac{M_2}{2} \right\rceil = \left\lceil \frac{m_2 + q - 1}{2} \right\rceil, \tag{C.58}$$

with q having its usual definition as the maximum of the reaction orders in the chemical system. As described in Section 4.3.1 and Section 4.3.4, each vector $\tilde{\boldsymbol{\mu}}(T)$ and $\tilde{\mathbf{z}}^{(\rho)}$ for $\rho \in \mathcal{R}$ contains only moments up to order $2n_2$. The objective value associated with the feasible point $(\tilde{\boldsymbol{\mu}}(T), (\tilde{\mathbf{z}}^{(\rho)})_{\rho \in \mathcal{R}})$ is $\tilde{\mu}_{\mathbf{e}_i}(T)$.

Next, we want to show that we can use $(\tilde{\boldsymbol{\mu}}(T), (\tilde{\mathbf{z}}^{(\rho)})_{\rho \in \mathcal{R}})$ to construct a point $(\hat{\boldsymbol{\mu}}(T), (\hat{\mathbf{z}}^{(\rho)})_{\rho \in \mathcal{R}})$ which is feasible for SDP₁ and has the same objective value of $\tilde{\mu}_{\mathbf{e}_i}(T)$. To be feasible, it must satisfy

$$\hat{\mu}_0(T) = 1 \tag{C.59}$$

$$\hat{\boldsymbol{\mu}}(T) \in C_{n_1}(\boldsymbol{\alpha}, \boldsymbol{\beta}) \tag{C.60}$$

$$\hat{\mathbf{z}}^{(\rho)} \in C_{n_1}(\boldsymbol{\alpha}, \boldsymbol{\beta}), \quad \forall \rho \in \mathcal{R}$$
 (C.61)

$$\hat{\boldsymbol{\mu}}_{L}(T) - e^{\rho T} \boldsymbol{\mu}_{L}(0) = (\mathbf{A}_{L}^{(1)} - \rho \mathbf{I}) \hat{\mathbf{z}}_{L}^{(\rho)} + \mathbf{A}_{H}^{(1)} \hat{\mathbf{z}}_{H}^{(\rho)}, \quad \forall \rho \in \mathcal{R}$$
(C.62)

$$\hat{z}_{\mathbf{0}}^{(\rho)} = \begin{cases} T & \text{if } \rho = 0, \\ \frac{e^{\rho T} - 1}{\rho} & \text{otherwise} \end{cases} \quad \forall \rho \in \mathcal{R},$$
(C.63)

where the matrices $\mathbf{A}_{L}^{(1)}$ and $\mathbf{A}_{H}^{(1)}$ contain rows corresponding to all moments up to

order m_1 , and where

$$n_1 = \left\lceil \frac{M_1}{2} \right\rceil = \left\lceil \frac{m_1 + q - 1}{2} \right\rceil. \tag{C.64}$$

We have two cases to consider:

• <u>Case 1</u>: $m_2 + q - 1$ is even.

In this case, $m_1 + q - 1 = (m_2 - 1) + q - 1 = (m_2 + q - 1) - 1$ is odd, and

$$n_1 = \left\lceil \frac{M_1}{2} \right\rceil = \left\lceil \frac{m_1 + q - 1}{2} \right\rceil = \left\lceil \frac{(m_2 + q - 1) - 1}{2} \right\rceil = n_2$$
(C.65)

It follows that $(\hat{\boldsymbol{\mu}}(T), (\hat{\mathbf{z}}^{(\rho)})_{\rho \in \mathcal{R}}) \equiv (\tilde{\boldsymbol{\mu}}(T), (\tilde{\mathbf{z}}^{(\rho)})_{\rho \in \mathcal{R}})$ has the appropriate dimension to be considered as a decision vector for SDP₁. In fact, it is also feasible in that it satisfies Conditions (C.59) - (C.63). Since $(\hat{\boldsymbol{\mu}}(T), (\hat{\mathbf{z}}^{(\rho)})_{\rho \in \mathcal{R}}) = (\tilde{\boldsymbol{\mu}}(T), (\tilde{\mathbf{z}}^{(\rho)})_{\rho \in \mathcal{R}})$ and $n_1 = n_2$, the satisfaction of all constraints except Condition (C.62) is trivial. Condition (C.62) is satisfied because Condition (C.62) enforces only a subset of the linear equations enforced by Condition (C.56) (i.e., the rows of the matrix $[\mathbf{A}_L^{(1)}\mathbf{A}_H^{(1)}]$ are a subset of those of the matrix $[\mathbf{A}_L^{(2)}\mathbf{A}_H^{(2)}]$). Finally, since $\hat{\mu}_{\mathbf{e}_i}(T) = \tilde{\mu}_{\mathbf{e}_i}(T)$, the objective value of $(\hat{\boldsymbol{\mu}}(T), (\hat{\mathbf{z}}^{(\rho)})_{\rho \in \mathcal{R}})$ in SDP₁ is identical to the objective value of $(\tilde{\boldsymbol{\mu}}(T), (\tilde{\mathbf{z}}^{(\rho)})_{\rho \in \mathcal{R}})$ in SDP₂.

• <u>Case 2</u>: $m_2 + q - 1$ is odd.

In this case, $m_1 + q - 1 = (m_2 - 1) + q - 1 = (m_2 + q - 1) - 1$ is even and

$$n_1 = \left\lceil \frac{M_1}{2} \right\rceil = \left\lceil \frac{m_1 + q - 1}{2} \right\rceil = \left\lceil \frac{(m_2 + q - 1) - 1}{2} \right\rceil = n_2 - 1$$
(C.66)

A decision vector $(\hat{\boldsymbol{\mu}}(T), (\hat{\mathbf{z}}^{(\rho)})_{\rho \in \mathcal{R}})$ for SDP₁ can only contain those moments up to order $2n_1 = 2(n_2 - 1) = 2n_2 - 2$. So, let $(\hat{\boldsymbol{\mu}}(T), (\hat{\mathbf{z}}^{(\rho)})_{\rho \in \mathcal{R}})$ be a truncation of $(\tilde{\boldsymbol{\mu}}(T), (\tilde{\mathbf{z}}^{(\rho)})_{\rho \in \mathcal{R}})$, in which only those moments up to order $2n_1 = 2n_2 - 2$ appear. This $(\hat{\boldsymbol{\mu}}(T), (\hat{\mathbf{z}}^{(\rho)})_{\rho \in \mathcal{R}})$ is feasible for SDP₁, though feasibility is a little harder to show than it was for Case 1. Let us consider Conditions (C.59) - (C.63) one at a time. First, since $\hat{\mu}_0(T) = \tilde{\mu}_0(T)$ and $\hat{z}_0^{(\rho)} = \tilde{z}_0^{(\rho)}$, Conditions (C.59) and (C.63) are satisfied trivially. To show that Condition (C.54) is satisfied, we recall that $\tilde{\boldsymbol{\mu}}(T) \in C_{n_2}(\boldsymbol{\alpha}, \boldsymbol{\beta})$ is equivalent to

$$\mathbf{M}_{n_2}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}(T)) \succeq \mathbf{0}, \tag{C.67}$$

$$\mathbf{M}_{n_2-1}^{\mathbf{e}_j}(\tilde{\boldsymbol{\mu}}(T)) \succeq \mathbf{0}, \quad \forall j \in \{1, \dots, \hat{N}\},$$
(C.68)

$$\alpha_k \mathbf{M}_{n_2-1}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}(T)) - \sum_{j=1}^{\hat{N}} \beta_{k,j} \mathbf{M}_{n_2-1}^{\mathbf{e}_j}(\tilde{\boldsymbol{\mu}}(T)) \succeq \mathbf{0}, \quad \forall k \in \{1, \dots, L\}.$$
(C.69)

The matrix $\mathbf{M}_{n_1}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}(T)) = \mathbf{M}_{n_2-1}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}(T))$ is a principal submatrix of $\mathbf{M}_{n_2}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}(T))$. Thus, $\mathbf{M}_{n_2}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}(T)) \succeq \mathbf{0}$ implies $\mathbf{M}_{n_1}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}(T)) \succeq \mathbf{0}$. Furthermore, since $\mathbf{M}_{n_1}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}(T))$ contains only moments up to order $2n_1$, we have $\mathbf{M}_{n_1}^{\mathbf{0}}(\tilde{\boldsymbol{\mu}}(T)) = \mathbf{M}_{n_1}^{\mathbf{0}}(\hat{\boldsymbol{\mu}}(T))$. It follows that

$$\mathbf{M}_{n_1}^{\mathbf{0}}(\hat{\boldsymbol{\mu}}(T)) \succeq \mathbf{0} \tag{C.70}$$

By essentially the same argument, we can conclude that

$$\mathbf{M}_{n_1-1}^{\mathbf{e}_j}(\hat{\boldsymbol{\mu}}(T)) \succeq \mathbf{0}, \quad \forall j \in \{1, \dots, \hat{N}\},$$
(C.71)

and

$$\alpha_k \mathbf{M}_{n_1-1}^{\mathbf{0}}(\hat{\boldsymbol{\mu}}(T)) - \sum_{j=1}^{\hat{N}} \beta_{k,j} \mathbf{M}_{n_1-1}^{\mathbf{e}_j}(\hat{\boldsymbol{\mu}}(T)) \succeq \mathbf{0}, \quad \forall k \in \{1, \dots, L\}, \qquad (C.72)$$

as well. LMIs (C.70) - (C.72) together are equivalent to the statement that $\hat{\mu}(T) \in C_{n_1}(\alpha, \beta)$. Thus, we see that Condition (C.60) is satisfied. By the same logic, we can show that Condition (C.61) is satisfied. All that remains, then, is to show the Condition (C.62) is satisfied. As in our treatment of Case 1, we again have that Condition (C.62) enforces only a subset of the linear equations enforced by Condition (C.56) (i.e., the rows of the matrix $[\mathbf{A}_L^{(1)}\mathbf{A}_H^{(1)}]$ are a subset of those of the matrix $[\mathbf{A}_L^{(2)}\mathbf{A}_H^{(2)}]$. In this Case, however, we also have the important fact that the elements of $(\tilde{\boldsymbol{\mu}}(T), (\tilde{\mathbf{z}}^{(\rho)})_{\rho \in \mathcal{R}})$ which don't appear in

 $(\hat{\boldsymbol{\mu}}(T), (\hat{\mathbf{z}}^{(\rho)})_{\rho \in \mathcal{R}})$ also do not appear in the linear equations of Conditon (C.62). This is because the linear equations of Conditon (C.62) are derived from the expressions for the time derivatives of the moments up to order m_1 , which involve only moments up to order

$$m_1 + q - 1 = 2\left(\frac{m_1 + q - 1}{2}\right) \le 2\left\lceil\frac{m_1 + q - 1}{2}\right\rceil = 2n_1.$$
 (C.73)

Thus, Condition (C.62) is satisfied, and we have shown that $(\hat{\boldsymbol{\mu}}(T), (\hat{\mathbf{z}}^{(\rho)})_{\rho \in \mathcal{R}})$ is feasible for SDP₁. Finally, since $\hat{\mu}_{\mathbf{e}_i}(T) = \tilde{\mu}_{\mathbf{e}_i}(T)$, the objective value of $(\hat{\boldsymbol{\mu}}(T), (\hat{\mathbf{z}}^{(\rho)})_{\rho \in \mathcal{R}})$ in SDP₁ is identical to the objective value of $(\tilde{\boldsymbol{\mu}}(T), (\tilde{\mathbf{z}}^{(\rho)})_{\rho \in \mathcal{R}})$ in SDP₂.

This completes the proof.

While the above proof was concerned with the upper bound on the mean molecular count of independent species i, specifically, these details are not essential to the argument. The same basic argument can be used to show that

- lower bounds on mean molecular counts,
- bounds on mean molecular counts for dependent species, and
- upper bounds on variances

all vary monotonically with increasing m.

Appendix D

Supplemental Material for Chapter 5

D.1 Introduction

This appendix contains material to supplement Chapter 5.

D.2 A Moment Enumeration Scheme

To implement the methods described in Chapter 5, it is necessary to be able to list systematically all moments $y_{\mathbf{j}}$ up to a specified order. In particular, this is necessary for constructing the moment vector \mathbf{y} , the matrices $\mathbf{A}_{L}^{\mathcal{E}}$ and $\mathbf{A}_{H}^{\mathcal{E}}$, and the matrices $\mathbf{M}_{n}^{\mathbf{0}}(\mathbf{y})$, $\mathbf{M}_{n-1}^{\mathbf{0}}(\mathbf{y})$, and $\mathbf{M}_{n-1}^{\mathbf{e}_{r}}(\mathbf{y})$. In Appendix B, we showed how this could be done in the context of the multi-indices $\mathbf{j} \in \mathbb{N}^{\hat{N}}$ and the moments $\boldsymbol{\mu}$. The generalization of these ideas to the multi-indices $\mathbf{j} \in \mathbb{N}^{R}$ and the moments \mathbf{y} is straightforward. While we defined the ordering relation " \leq_{gd} " for multi-indices $\mathbf{j} \in \mathbb{N}^{\hat{N}}$, the generalization to multi-indices of arbitrary dimension is straightforward. In particular, the ordering relation and the above algorithms can also be applied to the multi-indices $\mathbf{j} \in \mathbb{N}^{R}$ used to index the moments \mathbf{y} .

D.3 Definition of $\mathbf{M}_{n}^{\mathbf{0}}(\mathbf{y})$, $\mathbf{M}_{n-1}^{\mathbf{0}}(\mathbf{y})$, and $\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\mathbf{y})$

The definitions of the matrices $\mathbf{M}_{n}^{\mathbf{0}}(\mathbf{y})$, $\mathbf{M}_{n-1}^{\mathbf{0}}(\mathbf{y})$, and $\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\mathbf{y})$ exactly parallel the definitions of the matrices $\mathbf{M}_{n}^{\mathbf{0}}(\boldsymbol{\mu})$, $\mathbf{M}_{n-1}^{\mathbf{0}}(\boldsymbol{\mu})$, and $\mathbf{M}_{n-1}^{\mathbf{e}_{j}}(\boldsymbol{\mu})$ given in Appendix B. The only difference is that each multi-index is now an element of \mathbb{N}^{R} instead of $\mathbb{N}^{\hat{N}}$. Accordingly, a matrix whose subscript is $k \in \mathbb{N}$ (e.g., $\mathbf{M}_{k}^{\mathbf{0}}(\mathbf{y})$) has $\binom{k+R}{k}$ rows and columns, because that is the number of multi-indices $\mathbf{j} \in \mathbb{N}^{R}$ up to order k.

D.4 Derivation of LMIs

While we do not explicitly derive any of the LMIs involving the moment vector \mathbf{y} , their derivations are directly analogous to those of the LMIs involving the moment vector $\boldsymbol{\mu}$, in Appendix B. We simply use \mathbf{y} in place of $\boldsymbol{\mu}$, $P_{\mathcal{E}}(\cdot, t)$ in place of $P(\cdot, t)$, and different inequalities describing the support of the distribution.

D.5 Construction of the C matrix

This section describes the construction of the C matrix, which is used to transform the time derivatives of the vector $\mathbf{y}_L(t)$ into the time derivatives of the vector $\boldsymbol{\mu}_L(t)$.

D.5.1 Linear Mapping Between Moments

First, setting aside time derivatives, we show that the moments $\boldsymbol{\mu}_L(t)$ can be written as a linear function of the moments $\mathbf{y}_L(t)$. Picking an arbitrary multi-index $\mathbf{j} \in \mathbb{N}^{\hat{N}}$, we have

$$\mu_{\mathbf{j}} = \sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}} \hat{\mathbf{x}}^{\mathbf{j}} \hat{P}(\hat{\mathbf{x}}, t),
= \sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}} \hat{\mathbf{x}}^{\mathbf{j}} \sum_{\boldsymbol{\epsilon}\in\mathcal{E}(\hat{\mathbf{x}})} P_{\mathcal{E}}(\boldsymbol{\epsilon}, t),
= \sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}} \sum_{\boldsymbol{\epsilon}\in\mathcal{E}(\hat{\mathbf{x}})} \hat{\mathbf{x}}^{\mathbf{j}} P_{\mathcal{E}}(\boldsymbol{\epsilon}, t),
= \sum_{\hat{\mathbf{x}}\in\hat{\mathcal{X}}} \sum_{\boldsymbol{\epsilon}\in\mathcal{E}(\hat{\mathbf{x}})} \left(\hat{\mathbf{x}}_{0} + \hat{\mathbf{S}}\boldsymbol{\epsilon} \right)^{\mathbf{j}} P_{\mathcal{E}}(\boldsymbol{\epsilon}, t),
= \sum_{\boldsymbol{\epsilon}\in\mathcal{E}} \left(\hat{\mathbf{x}}_{0} + \hat{\mathbf{S}}\boldsymbol{\epsilon} \right)^{\mathbf{j}} P_{\mathcal{E}}(\boldsymbol{\epsilon}, t),
= \sum_{\boldsymbol{\epsilon}\in\mathcal{E}} \prod_{n=1}^{\hat{N}} \left(\hat{x}_{0,n} + \sum_{r=1}^{R} \hat{s}_{n,r}\boldsymbol{\epsilon}_{r} \right)^{j_{n}} P_{\mathcal{E}}(\boldsymbol{\epsilon}, t).$$
(D.1)

Applying the multinomial formula, we have

$$\mu_{\mathbf{j}} = \sum_{\boldsymbol{\epsilon}\in\mathcal{E}} \prod_{n=1}^{\hat{N}} \left(\sum_{|\mathbf{i}|\leq j_n} \begin{pmatrix} j_n \\ j_n - |\mathbf{i}|, i_1, \dots, i_R \end{pmatrix} \hat{x}_{0,n}^{j_n - |\mathbf{i}|} \prod_{r=1}^R (\hat{s}_{n,r}\epsilon_r)^{i_r} \right) P_{\mathcal{E}}(\boldsymbol{\epsilon}, t),$$

$$= \sum_{\boldsymbol{\epsilon}\in\mathcal{E}} \prod_{n=1}^{\hat{N}} \left(\sum_{|\mathbf{i}|\leq j_n} \begin{pmatrix} j_n \\ j_n - |\mathbf{i}|, i_1, \dots, i_R \end{pmatrix} \hat{x}_{0,n}^{j_n - |\mathbf{i}|} \left(\prod_{r=1}^R \hat{s}_{n,r}^{i_r} \right) \boldsymbol{\epsilon}^{\mathbf{i}} \right) P_{\mathcal{E}}(\boldsymbol{\epsilon}, t).$$
(D.2)

The expression inside the outermost parentheses is a polynomial in $\boldsymbol{\epsilon}$ and can thus be written as concisely as $\sum_{|\mathbf{i}| \leq j_n} \phi_{n,j_n,\mathbf{i}} \boldsymbol{\epsilon}^{\mathbf{i}}$, where

$$\phi_{n,j_n,\mathbf{i}} \equiv \sum_{|\mathbf{i}| \le j_n} \binom{j_n}{j_n - |\mathbf{i}|, i_1, \dots, i_R} \hat{x}_{0,n}^{j_n - |\mathbf{i}|} \left(\prod_{r=1}^R \hat{s}_{n,r}^{i_r}\right)$$
(D.3)

is easily computable. Substituting this expression back into Equation (D.2), and introducing distinct multi-indices $\mathbf{i}(n) \in \mathbb{N}^R$ for each $n \in \{1, \dots, \hat{N}\}$, we obtain

$$\mu_{\mathbf{j}} = \sum_{\boldsymbol{\epsilon}\in\mathcal{E}} \prod_{n=1}^{\hat{N}} \left(\sum_{|\mathbf{i}|\leq j_n} \phi_{n,j_n,\mathbf{i}} \boldsymbol{\epsilon}^{\mathbf{i}} \right) P_{\mathcal{E}}(\boldsymbol{\epsilon}, t),$$

$$= \sum_{\boldsymbol{\epsilon}\in\mathcal{E}} \prod_{n=1}^{\hat{N}} \left(\sum_{|\mathbf{i}(n)|\leq j_n} \phi_{n,j_n,\mathbf{i}(n)} \boldsymbol{\epsilon}^{\mathbf{i}(n)} \right) P_{\mathcal{E}}(\boldsymbol{\epsilon}, t).$$
(D.4)

Finally, we exchange the order of the inner sum and product and simplify:

$$\mu_{\mathbf{j}} = \sum_{\boldsymbol{\epsilon}\in\mathcal{E}} \sum_{|\mathbf{i}(i)|\leq j_{1}} \cdots \sum_{|\mathbf{i}(\hat{N})|\leq j_{\hat{N}}} \left(\prod_{n=1}^{\hat{N}} \phi_{n,j_{n},\mathbf{i}(n)} \boldsymbol{\epsilon}^{\mathbf{i}(n)} \right) P_{\mathcal{E}}(\boldsymbol{\epsilon},t),$$

$$= \sum_{\boldsymbol{\epsilon}\in\mathcal{E}} \sum_{|\mathbf{i}(i)|\leq j_{1}} \cdots \sum_{|\mathbf{i}(\hat{N})|\leq j_{\hat{N}}} \left(\prod_{n=1}^{\hat{N}} \phi_{n,j_{n},\mathbf{i}(n)} \right) \boldsymbol{\epsilon}^{\mathbf{i}(1)+\dots+\mathbf{i}(\hat{N})} P_{\mathcal{E}}(\boldsymbol{\epsilon},t),$$

$$= \sum_{|\mathbf{i}(i)|\leq j_{1}} \cdots \sum_{|\mathbf{i}(\hat{N})|\leq j_{\hat{N}}} \left(\prod_{n=1}^{\hat{N}} \phi_{n,j_{n},\mathbf{i}(n)} \right) \sum_{\boldsymbol{\epsilon}\in\mathcal{E}} \boldsymbol{\epsilon}^{\mathbf{i}(1)+\dots+\mathbf{i}(\hat{N})} P_{\mathcal{E}}(\boldsymbol{\epsilon},t),$$

$$= \sum_{|\mathbf{i}(i)|\leq j_{1}} \cdots \sum_{|\mathbf{i}(\hat{N})|\leq j_{\hat{N}}} \left(\prod_{n=1}^{\hat{N}} \phi_{n,j_{n},\mathbf{i}(n)} \right) y_{\mathbf{i}(1)+\dots+\mathbf{i}(\hat{N})}(t).$$
(D.5)

The coefficients in the last expression can be calculated using \hat{N} nested "for" loops.

Equation (D.5) expresses $\mu_{\mathbf{j}}(t)$ as a linear combination of moments $\{y_{\mathbf{k}}\}_{|\mathbf{k}| \leq |\mathbf{j}|}$. This implies that we can express $\boldsymbol{\mu}_{L}(t)$, all moments up to arbitrary order m, as a linear combination of $\{y_{\mathbf{k}}\}_{|\mathbf{k}| \leq m}$. This linear relationship can be expressed concisely as

$$\boldsymbol{\mu}_L(t) = \mathbf{C} \mathbf{y}_L(t). \tag{D.6}$$

Let $C_{\mathbf{j},\mathbf{k}}$ be the element of \mathbf{C} in the row corresponding to $\mathbf{j} \in \mathbb{N}^{\hat{N}}$ and the column corresponding to $\mathbf{k} \in \mathbb{N}^{R}$. Equation (D.5) gives us a recipe for computing $C_{\mathbf{j},\mathbf{k}}$ for all $\mathbf{k} \in \mathbb{N}^{R}$. The coefficient corresponding to the multi-indices $\mathbf{i}(1), \ldots, \mathbf{i}(\hat{N})$ contributes to $C_{\mathbf{j}, \mathbf{i}(1)+\cdots+\mathbf{i}(\hat{N})}$. Note that several combinations of the multi-indices $\mathbf{i}(1), \ldots, \mathbf{i}(\hat{N})$ may give the same summed multi-index $\mathbf{i}(1) + \cdots + \mathbf{i}(\hat{N}) \equiv \mathbf{k}$. Thus, in general, it will be necessary to add together several coefficients to compute $C_{\mathbf{j},\mathbf{k}}$.

D.5.2 Linear Mapping Between Time Derivatives

We can obtain an expression for $d\mu_L/dt$ in terms of $d\mathbf{y}_L/dt$ by simply differentiating both sides of Equation (D.6) with respect to time:

$$\frac{d\boldsymbol{\mu}_L}{dt}(t) = \mathbf{C}\frac{d\mathbf{y}_L}{dt}(t). \tag{D.7}$$

D.5.3 Generalization

Note that we could also express moments of the full-dimensional distribution $P(\cdot, t)$ in terms of the moments $\mathbf{y}(t)$. The derivation would be essentially the same as the one we have provided above. One would just have to remove the "hats" wherever they appear.

D.6 Construction of the $A^{\mathcal{E}}$ matrix

To construct the $\mathbf{A}^{\mathcal{E}}$ matrix, we need to be able to express the time derivative of an arbitrary moment $y_{\mathbf{j}}$ as a linear function of finitely many other moments. The first step toward this goal is to expand the expression $dy_{\mathbf{j}}/dt$ using the definition of $y_{\mathbf{j}}$

$$\frac{dy_{\mathbf{j}}}{dt}(t) = \frac{d}{dt} \left(\sum_{\boldsymbol{\epsilon} \in \mathcal{E}} \boldsymbol{\epsilon}^{\mathbf{j}} P_{\mathcal{E}}(\boldsymbol{\epsilon}, t) \right), \\
= \sum_{\boldsymbol{\epsilon} \in \mathcal{E}} \boldsymbol{\epsilon}^{\mathbf{j}} \left(\frac{dP_{\mathcal{E}}}{dt}(\boldsymbol{\epsilon}, t) \right).$$
(D.8)

The expression in parentheses can then be expanded using the alternative version of the CME we introduced in Chapter 5.

$$\frac{dy_{\mathbf{j}}}{dt}(t) = \sum_{\boldsymbol{\epsilon}\in\mathcal{E}} \boldsymbol{\epsilon}^{\mathbf{j}} \left(\sum_{r=1}^{R} \left[P_{\mathcal{E}}(\boldsymbol{\epsilon} - \mathbf{e}_{r}, t) \alpha_{r}(\boldsymbol{\epsilon} - \mathbf{e}_{r}) - P_{\mathcal{E}}(\boldsymbol{\epsilon}, t) \alpha_{r}(\boldsymbol{\epsilon}) \right] \right),$$

$$= \sum_{r=1}^{R} \sum_{\boldsymbol{\epsilon}\in\mathcal{E}} \boldsymbol{\epsilon}^{\mathbf{j}} \left(P_{\mathcal{E}}(\boldsymbol{\epsilon} - \mathbf{e}_{r}, t) \alpha_{r}(\boldsymbol{\epsilon} - \mathbf{e}_{r}) - P_{\mathcal{E}}(\boldsymbol{\epsilon}, t) \alpha_{r}(\boldsymbol{\epsilon}) \right),$$

$$= \sum_{r=1}^{R} \left(\sum_{\boldsymbol{\epsilon}\in\mathcal{E}} \boldsymbol{\epsilon}^{\mathbf{j}} P_{\mathcal{E}}(\boldsymbol{\epsilon} - \mathbf{e}_{r}, t) \alpha_{r}(\boldsymbol{\epsilon} - \mathbf{e}_{r}) - \sum_{\boldsymbol{\epsilon}\in\mathcal{E}} \boldsymbol{\epsilon}^{\mathbf{j}} P_{\mathcal{E}}(\boldsymbol{\epsilon}, t) \alpha_{r}(\boldsymbol{\epsilon}) \right).$$
(D.9)

Next, we would like to replace each of the sums over \mathcal{E} with sums over \mathbb{Z}^R . For the second sum, it is relatively easy to see that

$$\sum_{\boldsymbol{\epsilon}\in\mathcal{E}}\boldsymbol{\epsilon}^{\mathbf{j}}P_{\mathcal{E}}(\boldsymbol{\epsilon},t)\alpha_{r}(\boldsymbol{\epsilon}) = \sum_{\boldsymbol{\epsilon}\in\mathbb{Z}^{R}}\boldsymbol{\epsilon}^{\mathbf{j}}P_{\mathcal{E}}(\boldsymbol{\epsilon},t)\alpha_{r}(\boldsymbol{\epsilon}), \qquad (D.10)$$

because $\mathcal{E} \subset \mathbb{Z}^R$, and for every $\boldsymbol{\epsilon} \in \mathbb{Z}^R$ such that $\boldsymbol{\epsilon} \notin \mathcal{E}$ we have $P_{\mathcal{E}}(\boldsymbol{\epsilon}, t) = 0$. The argument for the first sum is a little more complicated, so we will formalize it as a

claim with a proof.

Claim 12.

$$\sum_{\boldsymbol{\epsilon}\in\mathcal{E}}\boldsymbol{\epsilon}^{\mathbf{j}}P_{\mathcal{E}}(\boldsymbol{\epsilon}-\mathbf{e}_{r},t)\alpha_{r}(\boldsymbol{\epsilon}-\mathbf{e}_{r}) = \sum_{\boldsymbol{\epsilon}\in\mathbb{Z}^{R}}\boldsymbol{\epsilon}^{\mathbf{j}}P_{\mathcal{E}}(\boldsymbol{\epsilon}-\mathbf{e}_{r},t)\alpha_{r}(\boldsymbol{\epsilon}-\mathbf{e}_{r}), \quad (D.11)$$

where the only change is in the set from which the index of summation is drawn.

Proof. First, since $\mathcal{E} \subset \mathbb{Z}^R$, it follows that every term that appears on the left-hand sum also appears in the right-hand sum.

All that remains then is to show that all of the extra terms appearing in the right-hand sum evaluate to zero. In particular, we want to show that all the terms corresponding to $\boldsymbol{\epsilon} \in \mathbb{Z}^R$ such that $\boldsymbol{\epsilon} \notin \mathcal{E}$ evaluate to zero.

Pick an arbitrary $\boldsymbol{\epsilon} \in \mathbb{Z}^R$ such that $\boldsymbol{\epsilon} \notin \mathcal{E}$. Then, there are two cases to consider:

• <u>Case 1</u>: $\boldsymbol{\epsilon} - \mathbf{e}_r \notin \mathcal{E}$.

In this case, according to the definition of $P_{\mathcal{E}}(\cdot, t)$, we have $P_{\mathcal{E}}(\boldsymbol{\epsilon} - \mathbf{e}_r, t) = 0$. Thus, the term in the summation corresponding to $\boldsymbol{\epsilon}$ is zero.

• <u>Case 2</u>: $\epsilon - \mathbf{e}_r \in \mathcal{E}$.

In this case, it is quite possible that $P_{\mathcal{E}}(\boldsymbol{\epsilon} - \mathbf{e}_r, t) > 0$, so we will instead argue that $\alpha_r(\boldsymbol{\epsilon} - \mathbf{e}_r) = 0$.

In Chapter 5, in the discussion surrounding the discussion of the propensity functions, we observed that for any state $\mathbf{x} \in \mathcal{X}$ and any reaction $\rho \in \{1, \ldots, R\}$ such that $\mathbf{x} + \mathbf{s}_{\rho} \notin \mathcal{X}$, we have $a_{\rho}(\mathbf{x}) = 0$. We will apply this general statement to the particular state

$$\mathbf{x} \equiv \mathbf{x}_0 + \mathbf{S}(\boldsymbol{\epsilon} - \mathbf{e}_r) = \mathbf{x}_0 + \mathbf{S}\boldsymbol{\epsilon} - \mathbf{s}_r.$$
(D.12)

Since $\epsilon - \mathbf{e}_r \in \mathcal{E}$, it follows that $\mathbf{x} \in \mathcal{X}$. Furthermore, since $\epsilon \notin \mathcal{E}$, we have that $\mathbf{x}_0 + \mathbf{S} \boldsymbol{\epsilon} = \mathbf{x} + \mathbf{s}_r \notin \mathcal{X}$. From the general statement above, it follows that

 $a_r(\mathbf{x}) = 0$. Moreover, from the definition of $\alpha_r(\cdot)$, we have

$$a_r(\mathbf{x}) = a_r(\mathbf{x}_0 + \mathbf{S}(\boldsymbol{\epsilon} - \mathbf{e}_r)) = \alpha_r(\boldsymbol{\epsilon} - \mathbf{e}_r).$$
(D.13)

Thus we have $\alpha_r(\boldsymbol{\epsilon} - \mathbf{e}_r) = 0$. It follows that the term in the summation corresponding to $\boldsymbol{\epsilon}$ is zero.

We have seen, in both of the two possible cases, that the term in the summation corresponding to $\boldsymbol{\epsilon}$ is zero. This was true for an arbitrary $\boldsymbol{\epsilon} \in \mathbb{Z}^R$ such that $\boldsymbol{\epsilon} \notin \boldsymbol{\mathcal{E}}$. Thus, it must be true for all $\boldsymbol{\epsilon} \in \mathbb{Z}^R$ such that $\boldsymbol{\epsilon} \notin \boldsymbol{\mathcal{E}}$. This completes the proof. \Box

Using Equations (D.10) and (D.11), we can write

$$\frac{dy_{\mathbf{j}}}{dt}(t) = \sum_{r=1}^{R} \left(\sum_{\boldsymbol{\epsilon} \in \mathbb{Z}^{R}} \boldsymbol{\epsilon}^{\mathbf{j}} P_{\mathcal{E}}(\boldsymbol{\epsilon} - \mathbf{e}_{r}, t) \alpha_{r}(\boldsymbol{\epsilon} - \mathbf{e}_{r}) - \sum_{\boldsymbol{\epsilon} \in \mathbb{Z}^{R}} \boldsymbol{\epsilon}^{\mathbf{j}} P_{\mathcal{E}}(\boldsymbol{\epsilon}, t) \alpha_{r}(\boldsymbol{\epsilon}) \right).$$
(D.14)

Then, shifting the index of the first sum, we can rewrite this as

$$\frac{dy_{\mathbf{j}}}{dt}(t) = \sum_{r=1}^{R} \left(\sum_{\boldsymbol{\epsilon} - \mathbf{e}_r \in \mathbb{Z}^R} \boldsymbol{\epsilon}^{\mathbf{j}} P_{\mathcal{E}}(\boldsymbol{\epsilon} - \mathbf{e}_r, t) \alpha_r(\boldsymbol{\epsilon} - \mathbf{e}_r) - \sum_{\boldsymbol{\epsilon} \in \mathbb{Z}^R} \boldsymbol{\epsilon}^{\mathbf{j}} P_{\mathcal{E}}(\boldsymbol{\epsilon}, t) \alpha_r(\boldsymbol{\epsilon}) \right),$$

$$= \sum_{r=1}^{R} \left(\sum_{\boldsymbol{\epsilon} \in \mathbb{Z}^R} (\boldsymbol{\epsilon} + \mathbf{e}_r)^{\mathbf{j}} P_{\mathcal{E}}(\boldsymbol{\epsilon}, t) \alpha_r(\boldsymbol{\epsilon}) - \sum_{\boldsymbol{\epsilon} \in \mathbb{Z}^R} \boldsymbol{\epsilon}^{\mathbf{j}} P_{\mathcal{E}}(\boldsymbol{\epsilon}, t) \alpha_r(\boldsymbol{\epsilon}) \right),$$

$$= \sum_{r=1}^{R} \sum_{\boldsymbol{\epsilon} \in \mathbb{Z}^R} \left((\boldsymbol{\epsilon} + \mathbf{e}_r)^{\mathbf{j}} - \boldsymbol{\epsilon}^{\mathbf{j}} \right) P_{\mathcal{E}}(\boldsymbol{\epsilon}, t) \alpha_r(\boldsymbol{\epsilon}),$$

$$= \sum_{r=1}^{R} \sum_{\boldsymbol{\epsilon} \in \mathcal{E}} \left((\boldsymbol{\epsilon} + \mathbf{e}_r)^{\mathbf{j}} - \boldsymbol{\epsilon}^{\mathbf{j}} \right) P_{\mathcal{E}}(\boldsymbol{\epsilon}, t) \alpha_r(\boldsymbol{\epsilon}).$$
(D.15)

Writing the monomials explicitly as products, this becomes

$$\frac{dy_{\mathbf{j}}}{dt}(t) = \sum_{r=1}^{R} \sum_{\boldsymbol{\epsilon} \in \mathcal{E}} \left(\prod_{\rho=1}^{R} (\epsilon_{\rho} + \delta_{r=\rho})^{j_{\rho}} - \prod_{\rho=1}^{R} (\epsilon_{\rho})^{j_{\rho}} \right) P_{\mathcal{E}}(\boldsymbol{\epsilon}, t) \alpha_{r}(\boldsymbol{\epsilon}),$$

$$= \sum_{r=1}^{R} \sum_{\boldsymbol{\epsilon} \in \mathcal{E}} \left(\prod_{\substack{\rho=1\\ \rho \neq r}}^{R} (\epsilon_{\rho})^{j_{\rho}} \right) \left((\epsilon_{r} + 1)^{j_{r}} - (\epsilon_{r})^{j_{r}} \right) P_{\mathcal{E}}(\boldsymbol{\epsilon}, t) \alpha_{r}(\boldsymbol{\epsilon}).$$
(D.16)

Applying the binomial formula gives

$$\frac{dy_{\mathbf{j}}}{dt}(t) = \sum_{r=1}^{R} \sum_{\epsilon \in \mathcal{E}} \left(\prod_{\substack{\rho=1\\\rho \neq r}}^{R} (\epsilon_{\rho})^{j_{\rho}} \right) \left(\sum_{k=0}^{j_{r}} {j_{r} \choose k} (\epsilon_{r})^{j_{r}-k} - (\epsilon_{r})^{j_{r}} \right) P_{\mathcal{E}}(\boldsymbol{\epsilon}, t) \alpha_{r}(\boldsymbol{\epsilon}), \\
= \sum_{r=1}^{R} \sum_{\epsilon \in \mathcal{E}} \left(\prod_{\substack{\rho=1\\\rho \neq r}}^{R} (\epsilon_{\rho})^{j_{\rho}} \right) \left(\sum_{k=1}^{j_{r}} {j_{r} \choose k} (\epsilon_{r})^{j_{r}-k} \right) P_{\mathcal{E}}(\boldsymbol{\epsilon}, t) \alpha_{r}(\boldsymbol{\epsilon}), \\
= \sum_{r=1}^{R} \sum_{\epsilon \in \mathcal{E}} \boldsymbol{\epsilon}^{\mathbf{j}-j_{r}\mathbf{e}_{r}} \left(\sum_{k=1}^{j_{r}} {j_{r} \choose k} (\epsilon_{r})^{j_{r}-k} \right) P_{\mathcal{E}}(\boldsymbol{\epsilon}, t) \alpha_{r}(\boldsymbol{\epsilon}), \\
= \sum_{r=1}^{R} \sum_{\epsilon \in \mathcal{E}} \left(\sum_{k=1}^{j_{r}} {j_{r} \choose k} \boldsymbol{\epsilon}^{\mathbf{j}-j_{r}\mathbf{e}_{r}}(\epsilon_{r})^{j_{r}-k} \right) P_{\mathcal{E}}(\boldsymbol{\epsilon}, t) \alpha_{r}(\boldsymbol{\epsilon}), \\
= \sum_{r=1}^{R} \sum_{\epsilon \in \mathcal{E}} \left(\sum_{k=1}^{j_{r}} {j_{r} \choose k} \boldsymbol{\epsilon}^{\mathbf{j}-j_{r}\mathbf{e}_{r}}(\epsilon_{r})^{j_{r}-k} \right) P_{\mathcal{E}}(\boldsymbol{\epsilon}, t) \alpha_{r}(\boldsymbol{\epsilon}), \\
= \sum_{r=1}^{R} \sum_{\epsilon \in \mathcal{E}} \sum_{k=1}^{j_{r}} {j_{r} \choose k} \boldsymbol{\epsilon}^{\mathbf{j}-k\mathbf{e}_{r}} P_{\mathcal{E}}(\boldsymbol{\epsilon}, t) \alpha_{r}(\boldsymbol{\epsilon}), \\
= \sum_{r=1}^{R} \sum_{\epsilon \in \mathcal{E}} \sum_{k=1}^{j_{r}} {j_{r} \choose k} \boldsymbol{\epsilon}^{\mathbf{j}-k\mathbf{e}_{r}} P_{\mathcal{E}}(\boldsymbol{\epsilon}, t) \alpha_{r}(\boldsymbol{\epsilon}), \\
= \sum_{r=1}^{R} \sum_{\epsilon \in \mathcal{E}} \sum_{k=1}^{j_{r}} {j_{r} \choose k} \boldsymbol{\epsilon}^{\mathbf{j}-k\mathbf{e}_{r}} \Omega_{r}(\boldsymbol{\epsilon}) P_{\mathcal{E}}(\boldsymbol{\epsilon}, t).$$

Now, for each $r \in \{1, ..., R\}$, the function $\alpha_r(\boldsymbol{\epsilon}) = a_r(\mathbf{x}_0 + \mathbf{S}\boldsymbol{\epsilon})$ will be a polynomial with respect to $\boldsymbol{\epsilon}$ (see the definition of the propensity function given in Chapter 5). Moreover, if reaction r involves a collision between $q_r \in \mathbb{N}$ molecules, then the polynomial will be order q_r . This polynomial can be written as

$$\alpha_r(\boldsymbol{\epsilon}) \equiv \sum_{|\mathbf{i}| \le q_r} a_{r,\mathbf{i}} \boldsymbol{\epsilon}^{\mathbf{i}},\tag{D.18}$$

where the coefficients $\{a_{r,\mathbf{i}}\}_{|\mathbf{i}| \leq q_r}$ depend on \mathbf{x}_0 and \mathbf{S} . As we explain in a later section, these coefficients can be expressed in terms of Stirling numbers of the first kind. For now, though, suffice it to say that they can be computed.

Substituting this Equation (D.18) into the last line of Equation (D.17), we obtain

$$\frac{dy_{\mathbf{j}}}{dt}(t) = \sum_{r=1}^{R} \sum_{\boldsymbol{\epsilon} \in \mathcal{E}} \sum_{k=1}^{j_r} {j_r \choose k} \boldsymbol{\epsilon}^{\mathbf{j}-k\mathbf{e}_r} \left(\sum_{|\mathbf{i}| \le q_r} a_{r,\mathbf{i}} \boldsymbol{\epsilon}^{\mathbf{i}}\right) P_{\mathcal{E}}(\boldsymbol{\epsilon}, t),$$

$$= \sum_{r=1}^{R} \sum_{\boldsymbol{\epsilon} \in \mathcal{E}} \sum_{k=1}^{j_r} \sum_{|\mathbf{i}| \le q_r} a_{r,\mathbf{i}} {j_r \choose k} \boldsymbol{\epsilon}^{\mathbf{j}-k\mathbf{e}_r} \boldsymbol{\epsilon}^{\mathbf{i}} P_{\mathcal{E}}(\boldsymbol{\epsilon}, t),$$

$$= \sum_{r=1}^{R} \sum_{\boldsymbol{\epsilon} \in \mathcal{E}} \sum_{k=1}^{j_r} \sum_{|\mathbf{i}| \le q_r} a_{r,\mathbf{i}} {j_r \choose k} \boldsymbol{\epsilon}^{\mathbf{j}-k\mathbf{e}_r+\mathbf{i}} P_{\mathcal{E}}(\boldsymbol{\epsilon}, t).$$
(D.19)

Finally, rearranging, we obtain

$$\frac{dy_{\mathbf{j}}}{dt}(t) = \sum_{r=1}^{R} \sum_{k=1}^{j_{r}} \sum_{|\mathbf{i}| \le q_{r}} a_{r,\mathbf{i}} {j_{r} \choose k} \sum_{\boldsymbol{\epsilon} \in \mathcal{E}} \boldsymbol{\epsilon}^{\mathbf{j}-k\mathbf{e}_{r}+\mathbf{i}} P_{\mathcal{E}}(\boldsymbol{\epsilon}, t),$$

$$= \sum_{r=1}^{R} \sum_{k=1}^{j_{r}} \sum_{|\mathbf{i}| \le q_{r}} a_{r,\mathbf{i}} {j_{r} \choose k} y_{\mathbf{j}-k\mathbf{e}_{r}+\mathbf{i}},$$
(D.20)

which achieves our goal of expressing dy_j/dt as a linear combination of finitely many elements of **y**.

The matrix **A** can be constructed one row at a time, where each row corresponds to some **j** such that $|\mathbf{j}| \leq m$. (The multi-indices can be placed in "graded descending order", as described in Section B.0.1.) The entire matrix is initialized with zeros. Then, for each **j**, we iterate over three nested "for" loops, one for each of the finite summations in the last line of Equation (D.20). For each value of r, k, and **i**, we evaluate the coefficient $a_{r,\mathbf{i}} {j_r \choose k}$ associated with the moment $y_{\mathbf{j}-k\mathbf{e}_r+\mathbf{i}}$. We then add this coefficient to the current value in the element of **A** whose row is indexed by **j** and whose column is indexed by $\mathbf{j} - k\mathbf{e}_r + \mathbf{i}$. It is important that we add the coefficient to the current value and not overwrite it, because different indices r, k, and **i** in later iterations might lead to a sum $\mathbf{j} - k\mathbf{e}_r + \mathbf{i}$ which indexes the same column.

D.6.1 Computing the $a_{r,i}$ Coefficients

This section explains how to compute the coefficients $\{a_{r,i}\}_{|i| \leq q_r}$ in the polynomial

$$\sum_{|\mathbf{i}| \le q_r} a_{r,\mathbf{i}} \boldsymbol{\epsilon}^{\mathbf{i}} \equiv \alpha_r(\boldsymbol{\epsilon}) = a_r(\mathbf{x}_0 + \mathbf{S}\boldsymbol{\epsilon}) = c_r \prod_{j=1}^N \begin{pmatrix} x_{0,j} + \sum_{\rho=1}^R s_{j,\rho} \boldsymbol{\epsilon}_\rho \\ \gamma_{j,r} \end{pmatrix}.$$
 (D.21)

The key algebraic fact that allows us to compute these coefficients is that

$$\binom{t}{n} = \sum_{i=0}^{n} s(n,i) \frac{t^i}{n!},$$
(D.22)

where s(n, i) is a Stirling number of the first kind, defined by the recurrence relation

$$s(n+1,i) = -ns(n,i) + s(n,i-1),$$
(D.23)

with the initial conditions

$$s(0,0) = 1,$$
 (D.24)

and

$$s(n,0) = s(0,i) = 0,$$
 (D.25)

if n > 0. See Wikipedia's page of "Binomial Coefficients" for details.

Applying this formula to Equation (D.21), we obtain

$$\sum_{|\mathbf{i}| \le q_r} a_{r,\mathbf{i}} \boldsymbol{\epsilon}^{\mathbf{i}} = c_r \prod_{j=1}^N \left(\sum_{f=0}^{\gamma_{j,r}} \frac{s(\gamma_{j,r}, f)}{\gamma_{j,r}!} \left(x_{0,j} + \sum_{\rho=1}^R s_{j,\rho} \boldsymbol{\epsilon}_\rho \right)^f \right).$$
(D.26)

Applying the multinomial theorem we obtain

$$\sum_{|\mathbf{i}| \le q_r} a_{r,\mathbf{i}} \boldsymbol{\epsilon}^{\mathbf{i}} = c_r \prod_{j=1}^N \left(\sum_{f=0}^{\gamma_{j,r}} \frac{s(\gamma_{j,r}, f)}{\gamma_{j,r}!} \left(\sum_{|\mathbf{k}| \le f} \binom{f}{f - |\mathbf{k}|, k_1, \dots, k_R} \right) x_{0,j}^{f-|\mathbf{k}|} \prod_{\rho=1}^R (s_{j,\rho} \epsilon_{\rho})^{k_{\rho}} \right) \right),$$

$$= c_r \prod_{j=1}^N \left(\sum_{f=0}^{\gamma_{j,r}} \frac{s(\gamma_{j,r}, f)}{\gamma_{j,r}!} \left(\sum_{|\mathbf{k}| \le f} \binom{f}{f - |\mathbf{k}|, k_1, \dots, k_R} \right) x_{0,j}^{f-|\mathbf{k}|} \left(\prod_{\rho=1}^R (s_{j,\rho})^{k_{\rho}} \right) \boldsymbol{\epsilon}^{\mathbf{k}} \right) \right),$$

$$= c_r \prod_{j=1}^N \left(\sum_{f=0}^{\gamma_{j,r}} \sum_{|\mathbf{k}| \le f} \frac{s(\gamma_{j,r}, f)}{\gamma_{j,r}!} \binom{f}{f - |\mathbf{k}|, k_1, \dots, k_R} x_{0,j}^{f-|\mathbf{k}|} \left(\prod_{\rho=1}^R (s_{j,\rho})^{k_{\rho}} \right) \boldsymbol{\epsilon}^{\mathbf{k}} \right).$$
(D.27)

Let us focus on the expression within the outermost set of parentheses. This is a polynomial of order at most $\gamma_{j,r}$. Thus, there exist coefficients $\{b_{r,j,\mathbf{n}}\}_{|\mathbf{n}| \leq \gamma_{j,r}}$ such
that

$$\sum_{f=0}^{\gamma_{j,r}} \sum_{|\mathbf{k}| \le f} \frac{s(\gamma_{j,r}, f)}{\gamma_{j,r}!} \binom{f}{f - |\mathbf{k}|, k_1, \dots, k_R} x_{0,j}^{f-|\mathbf{k}|} \left(\prod_{\rho=1}^R (s_{j,\rho})^{k_\rho}\right) \boldsymbol{\epsilon}^{\mathbf{k}} \equiv \sum_{|\mathbf{n}| \le \gamma_{j,r}} b_{r,j,\mathbf{n}} \boldsymbol{\epsilon}^{\mathbf{n}}.$$
(D.28)

We can compute these coefficients by iterating over the finite sums on the left-hand side of Equation (D.28). Thus, moving forward, we will assume they are known. We can then considerably simplify Equation (D.27) by writing things in terms of the $\{b_{r,j,\mathbf{n}}\}_{|\mathbf{n}| \leq \gamma_{j,r}}$:

$$\sum_{\mathbf{i}|\leq q_r} a_{r,\mathbf{i}} \boldsymbol{\epsilon}^{\mathbf{i}} = c_r \prod_{j=1}^N \left(\sum_{|\mathbf{n}|\leq \gamma_{j,r}} b_{r,j,\mathbf{n}} \boldsymbol{\epsilon}^{\mathbf{n}} \right).$$
(D.29)

To exchange the order of the product and sum, we introduce a distinct multi-index $\mathbf{n}^{(j)}$ for each $j \in \{1, \ldots, N\}$. We can then write things equivalently as

$$\sum_{|\mathbf{i}| \le q_r} a_{r,\mathbf{i}} \boldsymbol{\epsilon}^{\mathbf{i}} = c_r \prod_{j=1}^N \left(\sum_{|\mathbf{n}^{(j)}| \le \gamma_{j,r}} b_{r,j,\mathbf{n}^{(j)}} \boldsymbol{\epsilon}^{\mathbf{n}^{(j)}} \right),$$

$$= c_r \sum_{|\mathbf{n}^{(1)}| \le \gamma_{1,r}} \cdots \sum_{|\mathbf{n}^{(N)}| \le \gamma_{N,r}} \prod_{j=1}^N b_{r,j,\mathbf{n}^{(j)}} \boldsymbol{\epsilon}^{\mathbf{n}^{(j)}},$$

$$= c_r \sum_{|\mathbf{n}^{(1)}| \le \gamma_{1,r}} \cdots \sum_{|\mathbf{n}^{(N)}| \le \gamma_{N,r}} \left(\prod_{j=1}^N b_{r,j,\mathbf{n}^{(j)}} \right) \boldsymbol{\epsilon}^{\mathbf{n}^{(1)} + \dots + \mathbf{n}^{(N)}}.$$
 (D.30)

The coefficients $\{a_{r,\mathbf{i}}\}_{|\mathbf{i}| \leq q_r}$ can be computed by iterating over the finite sums in the last line of Equation (D.30).

D.7 Revisiting a Non-Pathological Example

As pointed out in Chapter 5, the bounds produced by the alternative formulation of the bounding method, are not guaranteed to be strictly better than the bounds produced by the original formulation. In particular, if the set $\bar{\mathcal{X}}$ does not outerapproximate the set $\hat{\mathcal{X}}$ in the sense described in Chapter 5, then the bounds produced by the two methods will be the same.

To demonstrate this, we revisit an example from one of our previous papers [15,

Section IV]:

$$A + B \xrightarrow{c_1} C \tag{D.31}$$

where $c_1 = 1 \text{ s}^{-1}$ and the initial molecular counts are $\mathbf{x}_0 = (x_{A0}, x_{B0}, x_{C0}) = (3, 4, 0)$. The stoichiometry matrix for this system is a single column:

$$\mathbf{S} = \begin{bmatrix} -1\\ -1\\ 1 \end{bmatrix}, \tag{D.32}$$

whose left null space is spanned by the rows of the matrix

$$\mathbf{B} = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix}. \tag{D.33}$$

Thus, the invariant values are $\mathbf{f} = \mathbf{B}\mathbf{x}_0 = (3, 4)$. Letting

$$\tilde{\mathbf{B}} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad \text{and} \quad \hat{\mathbf{B}} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \tag{D.34}$$

and picking species C to be considered independent, the equation defining $\bar{\mathcal{X}}$ simplifies to

$$\bar{\mathcal{X}} = \left\{ x_C \in \mathbb{R} : \begin{array}{c} x_C \ge 0\\ 3 - x_C \ge 0 \end{array} \right\}.$$
 (D.35)

Now, inspecting Reaction System (D.31) and the initial condition, we see that the full dimensional reachable states of this system are

$$\mathcal{X} = \left\{ \begin{bmatrix} 3\\4\\0 \end{bmatrix}, \begin{bmatrix} 2\\3\\1 \end{bmatrix}, \begin{bmatrix} 1\\2\\2 \end{bmatrix}, \begin{bmatrix} 0\\1\\3 \end{bmatrix} \right\}.$$
(D.36)

Treating species C as independent, the reduced states are

$$\hat{\mathcal{X}} = \{0, 1, 2, 3\}.$$
 (D.37)

In this case, we see that $\bar{\mathcal{X}}$ is the convex hull of $\hat{\mathcal{X}}$. There is no $\hat{\mathbf{x}} \in \mathbb{N}$ such that $\hat{\mathbf{x}} \in \bar{\mathcal{X}}$ but $\hat{\mathbf{x}} \notin \hat{\mathcal{X}}$. Accordingly, we would expect that the bounds produced by the alternative formulation of the bounding method to match those produced by the original formulation. This is, indeed, the case. In our previous paper [15, Section IV C], we calculated time-varying bounds on the mean of each species using the old dynamic bounding formulation, with m = 3, $\mathcal{R} = \{0, -2, -6\}$, and a range of t values from t = 0 to t = 4. Repeating the analysis with the new formulation, we obtain the time-varying mean bounds shown in Figure D-1. Comparing this figure with the top panel of Figure 3 in our previous paper [15, Section IV C], we see that the bounds are exactly the same.



Figure D-1: Time-varying bounds on the mean molecular counts of species A, B, and C for Reaction System (D.31). These bounds were calculated by solving the dynamic bounding SDP based on reaction extents (and its minimization counterpart) over a range of t values, using m = 3 and $\mathcal{R} = \{0, -2, -6\}$. Compare with the top panel of Figure 3 in our previous paper [15, Section IV C].

Bibliography

- [1] Amir Ali Ahmadi, Sanjeeb Dash, and Georgina Hall. Optimization over structured subsets of positive semidefinite matrices via column generation. *Discrete Optimization*, 24:129–151, 2017.
- MOSEK ApS. The MOSEK optimization toolbox for MATLAB manual. Version 7.1 (Revision 28)., 2015.
- [3] Dimitris Bertsimas and Constantine Caramanis. Bounds on linear pdes via semidefinite optimization. *Mathematical Programming*, 108(1):135–158, 2006.
- [4] Jerry W Blankenship and James E Falk. Infinitely constrained optimization problems. Journal of Optimization Theory and Applications, 19(2):261–281, 1976.
- [5] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. 2010.
- [6] Richard D. Braatz. Advanced control of crystallization processes. Annual Reviews in Control, 26(1):87–99, 2002.
- [7] Céline Casenave, Denis Dochain, Graciela Alvarez, Marcela Arellano, Hayat Benkhelifa, and Denis Leducq. Model identification and reduction for the control of an ice cream crystallization process. *Chemical Engineering Science*, 119:274– 287, 2014.
- [8] Giuseppi Cogoni and Patrick J Frawley. Particle size distribution reconstruction using a finite number of its moments through artificial neural networks: A practical application. Crystal Growth & Design, 15(1):239–246, 2014.
- [9] PH Constantino, M Vlysidis, P Smadbeck, and YN Kaznessis. Modeling stochasticity in biochemical reaction networks. *Journal of Physics D: Applied Physics*, 49(9):093001, 2016.
- [10] David Cox, John Little, and Donal O'shea. Ideals, varieties, and algorithms, volume 3. Springer, 1992.
- [11] L. G M de Souza, G. Janiga, V. John, and D. Thevenin. Reconstruction of a distribution from a finite number of moments with an adaptive spline-based algorithm. *Chemical Engineering Science*, 65(9):2741–2750, 2010.

- [12] Garrett R Dowdy and Paul I Barton. Using semidefinite programming to calculate bounds on particle size distributions. *Chemical Engineering Science*, 171:599–613, 2017.
- [13] Garrett R Dowdy and Paul I Barton. Using semidefinite programming to calculate bounds on stochastic chemical kinetic systems at steady state. In *Computer Aided Chemical Engineering*, volume 40, pages 2239–2244. Elsevier, 2017.
- [14] Garrett R Dowdy and Paul I Barton. Bounds on stochastic chemical kinetic systems at steady state. The Journal of Chemical Physics, 148(8):084106, 2018.
- [15] Garrett R Dowdy and Paul I Barton. Dynamic bounds on stochastic chemical kinetic systems using semidefinite programming. The Journal of Chemical Physics, 149(7):074103, 2018.
- [16] Garrett R Dowdy and Paul I Barton. Dynamic bounds on stochastic chemical kinetic systems using semidefinite programming. arXiv preprint arXiv:1802.04409, 2018.
- [17] Rong Fan, Daniele L. Marchisio, and Rodney O. Fox. Application of the direct quadrature method of moments to polydisperse gas-solid fluidized beds. *Powder Technology*, 139(1):7–20, 2004.
- [18] Giovanni Fantuzzi, David Goluskin, Deqing Huang, and Sergei I Chernyshenko. Bounds for deterministic and stochastic dynamical systems using sum-of-squares optimization. SIAM Journal on Applied Dynamical Systems, 15(4):1962–1988, 2016.
- [19] FDA. Quality by Design for ANDAs : An Example for Immediate-Release Dosage Forms, Module 3 Quality 3.2.P.2 Pharmaceutical Development, April 2012. www. fda.gov.
- [20] Magne Fjeld, OA Asbjørnsen, and Karl Johan Åström. Reaction invariants and their importance in the analysis of eigenvectors, state observability and controllability of the continuous stirred tank reactor. *Chemical Engineering Science*, 29(9):1917–1926, 1974.
- [21] Zachary Fox, Gregor Neuert, and Brian Munsky. Finite state projection based bounds to compare chemical master equation models using single-cell data. *The Journal of chemical physics*, 145(7):074101, 2016.
- [22] Pascal Gahinet, Arkadii Nemirovskii, Alan J Laub, and Mahmoud Chilali. The lmi control toolbox. In *Decision and Control*, 1994., Proceedings of the 33rd IEEE Conference on, volume 3, pages 2038–2041. IEEE, 1994.
- [23] Semyon Aranovich Gershgorin. Uber die abgrenzung der eigenwerte einer matrix. Proceedings of the Russian Academy of Sciences, 6(6):749–754, 1931.

- [24] Khem Raj Ghusinga, Cesar A Vargas-Garcia, Andrew Lamperski, and Abhyudai Singh. Exact lower and upper bounds on stationary moments in stochastic biochemical systems. *Physical Biology*, 14(4):04LT01, 2017.
- [25] Colin S Gillespie. Moment-closure approximations for mass-action models. IET Systems Biology, 3(1):52–58, 2009.
- [26] Daniel T Gillespie. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *Journal of computational physics*, 22(4):403–434, 1976.
- [27] Daniel T Gillespie. Stochastic simulation of chemical kinetics. Annu. Rev. Phys. Chem., 58:35–55, 2007.
- [28] M Grant and S Boyd. Cvx: Matlab software for disciplined convex programming, version 2.1. http://cvxr.com/cvx, March 2014.
- [29] Ankit Gupta and Mustafa Khammash. A generic state-space decomposition method for analyzing stochastic biomolecular reaction networks. arXiv preprint arXiv:1505.06594, 2015.
- [30] Ankit Gupta, Jan Mikelson, and Mustafa Khammash. A finite state projection algorithm for the stationary solution of the chemical master equation. *The Journal of Chemical Physics*, 147(15):154101, 2017.
- [31] Desmond J Higham. Modeling and simulating chemical reactions. *SIAM review*, 50(2):347–368, 2008.
- [32] HM. Hulburt and S. Katz. Some problems in particle technology: A statistical mechanical formulation. *Chemical Engineering Science*, 19:555–574, 1963.
- [33] Kieran Hutton, Niall Mitchell, and Patrick J. Frawley. Particle size distribution reconstruction: The moment surface method. *Powder Technology*, 222:8–14, 2012.
- [34] V John, I Angelov, AA Öncül, and D Thévenin. Techniques for the reconstruction of a distribution from a finite number of its moments. *Chemical Engineering Science*, 62(11):2890–2904, 2007.
- [35] Vladimir Kazeev, Mustafa Khammash, Michael Nip, and Christoph Schwab. Direct solution of the chemical master equation using quantized tensor trains. *PLoS computational biology*, 10(3):e1003359, 2014.
- [36] Yiannis N Kaznessis. Statistical thermodynamics and stochastic kinetics: An introduction for engineers. Cambridge University Press, 2011.
- [37] Hendrik Anthony Kramers. Brownian motion in a field of force and the diffusion model of chemical reactions. *Physica*, 7(4):284–304, 1940.

- [38] Kartik Krishnan and John E Mitchell. Cutting plane methods for semidefinite programming. 2002.
- [39] Juan Kuntz, Philipp Thomas, Guy-Bart Stan, and Mauricio Barahona. Rigorous bounds on the stationary distributions of the chemical master equation via mathematical programming. arXiv preprint arXiv:1702.05468, 2017.
- [40] Jean B Lasserre. Global optimization with polynomials and the problem of moments. SIAM Journal on optimization, 11(3):796–817, 2001.
- [41] Jean B Lasserre. Bounds on measures satisfying moment conditions. Annals of Applied Probability, 12:1114–1137, 2002.
- [42] Jean Bernard Lasserre. Moments, positive polynomials and their applications, volume 1. World Scientific, 2009.
- [43] Johan Löfberg. YALMIP: A toolbox for modeling and optimization in matlab. In *Proceedings of the CACSD Conference, Taipei, Taiwan,* 2004.
- [44] John R. Martin, Julian F. Johnson, and Anthony R. Cooper. Mechanical properties of polymers: The influence of molecular weight and molecular weight distribution. *Journal of Macromolecular Science*, Part C, 8(1):57–199, 1972.
- [45] Robert McGraw, Paul J DeMott, and Colin D O'Dowd. Sparse aerosol models beyond the quadrature method of moments. In AIP Conference Proceedings, volume 1527, pages 651–654. AIP, 2013.
- [46] Robert McGraw, Seth Nemesure, and Stephen E. Schwartz. Properties and evolution of aerosols with size distributions having identical moments. *Journal* of Aerosol Science, 29(7):761–772, 1998.
- [47] Donald A McQuarrie. Stochastic approach to chemical kinetics. *Journal of Applied Probability*, 4(3):413–478, 1967.
- [48] JE Moyal. Stochastic processes and statistical physics. Journal of the Royal Statistical Society. Series B (Methodological), 11(2):150–210, 1949.
- [49] Brian Munsky and Mustafa Khammash. The finite state projection algorithm for the solution of the chemical master equation. The Journal of chemical physics, 124(4):044104, 2006.
- [50] Allan Myerson. *Handbook of Industrial Crystallization*. Butterworth-Heinemann, Woburn, 2002.
- [51] Mohammad Naghnaeian and Domitilla Del Vecchio. Robust moment closure method for the chemical master equation. In *Control Technology and Applications* (CCTA), 2017 IEEE Conference on, pages 967–972. IEEE, 2017.

- [52] Clifford T. Ociardhá, Kieran W. Hutton, Niall A. Mitchell, and Patrick J. Frawley. Simultaneous parameter estimation and optimization of a seeded antisolvent crystallization. *Crystal Growth and Design*, 12(11):5247–5261, 2012.
- [53] Doraiswami Ramkrishna. *Population Balances*. Academic Press, San Diego, 2000.
- [54] AD Randolph and MA Larson. Theory of Particulate Processes. Academic Press Inc., New York, 1988.
- [55] Youcef Saad. Numerical methods for large eigenvalue problems. Manchester University Press, 1992.
- [56] Yuta Sakurai and Yutaka Hori. A convex approach to steady state moment analysis for stochastic chemical reactions. In *Decision and Control (CDC)*, 2017 *IEEE 56th Annual Conference on*, pages 1206–1211. IEEE, 2017.
- [57] Yuta Sakurai and Yutaka Hori. Optimization-based synthesis of stochastic biocircuits with statistical specifications. *Journal of The Royal Society Interface*, 15(138):20170709, 2018.
- [58] Kevin R Sanft, Sheng Wu, Min Roh, Jin Fu, Rone Kwei Lim, and Linda R Petzold. Stochkit2: software for discrete stochastic simulation of biochemical systems with events. *Bioinformatics*, 27(17):2457–2458, 2011.
- [59] David Schnoerr, Guido Sanguinetti, and Ramon Grima. Approximation and inference methods for stochastic biochemical kinetics—a tutorial review. *Journal* of Physics A: Mathematical and Theoretical, 50(9):093001, 2017.
- [60] Dan Shi, Prashant Mhaskar, Nael H El-Farra, and Panagiotis D Christofides. Predictive control of crystal size distribution in protein crystallization. Nanotechnology, 16(7):S562–74, 2005.
- [61] Sergei L Shmakov. Algorithm to calculate limiting cumulative particle size distribution functions from turbidimetric data. *Applied Optics*, 53(2):301–305, 2014.
- [62] James Alexander Shohat and Jacob David Tamarkin. The problem of moments. Number 1. American Mathematical Soc., 1943.
- [63] P Smadbeck and YN Kaznessis. Efficient moment matrix generation for arbitrary chemical networks. *Chemical engineering science*, 84:612–618, 2012.
- [64] Patrick Smadbeck and Yiannis N Kaznessis. A closure scheme for chemical master equations. Proceedings of the National Academy of Sciences, 110(35):14261– 14265, 2013.
- [65] Vassilios Sotiropoulos and Yiannis N Kaznessis. Analytical derivation of moment equations in stochastic chemical kinetics. *Chemical Engineering Science*, 66(3):268–277, 2011.

- [66] J.F. Sturm. Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. Optimization Methods and Software, 11-12:625-653, 1999. Version 1.05 available from http://fewcal.kub.nl/sturm.
- [67] Michael Eugene Taylor. *Measure theory and integration*. American Mathematical Soc., 2006.
- [68] Philipp Thomas, Arthur V Straube, and Ramon Grima. The slow-scale linear noise approximation: an accurate, reduced stochastic description of biochemical networks under timescale separation conditions. *BMC systems biology*, 6(1):39, 2012.
- [69] Nicolaas G van Kampen. The expansion of the master equation. Adv. Chem. Phys, 34:245–309, 1976.
- [70] Jeremy G VanAntwerp and Richard D Braatz. A tutorial on linear and bilinear matrix inequalities. *Journal of process control*, 10(4):363–385, 2000.
- [71] Lieven Vandenberghe and Stephen Boyd. Semidefinite programming. SIAM Review, 38(1):49–95, 1996.
- [72] Huy D Vo and Roger B Sidje. An adaptive solution to the chemical master equation using tensors. *The Journal of chemical physics*, 147(4):044102, 2017.
- [73] Warren H. White. Particle size distributions that cannot be distinguished by their integral moments. Journal of Colloid And Interface Science, 135(1):297– 299, 1990.
- [74] C. Wibowo, WC. Chang, and KM. Ng. Design of integrated crystallization systems. AIChE Journal, 47:2474–2492, 2001.
- [75] Allan D. Woodbury. A FORTRAN program to produce minimum relative entropy distributions. *Computers and Geosciences*, 30(1):131–138, 2004.
- [76] Makoto Yamashita, Katsuki Fujisawa, and Masakazu Kojima. Implementation and evaluation of sdpa 6.0 (semidefinite programming algorithm 6.0). Optimization Methods and Software, 18(4):491–505, 2003.