

Robust model invalidation for chemical reaction networks using generalized moments (extended version)

Theodore W. Grunberg¹ and Domitilla Del Vecchio²

Abstract—Many biomolecular systems can be described by chemical reaction networks. Determining which chemical reactions network models are inconsistent with observed data can be done via model invalidation. In this work, we formulate and solve a robust version of the model invalidation problem for the case where only measurements from the stationary distribution are available. This problem corresponds to determining if an observed distribution could be generated by the given chemical reaction network for some value of the parameters, plus a perturbation of bounded size with respect to total variation distance. The main technical tool we introduce to solve the problem is a set of generalized moments that make the problem amenable to an algorithmic solution.

I. INTRODUCTION

Many systems of interest in synthetic and systems biology can be modeled by stochastic chemical reaction networks. These models describe the evolution of the molecular counts of each species as a Markov chain over the non-negative lattice [3]. For many design problems in synthetic biology, one must select molecules that implement a specific set of chemical reactions from a selection of possible candidate molecules. Therefore, determining from experimental data whether a particular set of candidate molecules implements the desired model is highly relevant for design. One framework for deciding if a model is “correct” is *model invalidation*. In this framework, the experimental data is either consistent or inconsistent with a given model. If the data is inconsistent with the model, then the model is invalidated since we know that the model is incorrect in the sense that it could not have generated the data. On the other hand, if the data is consistent with the model, the model may or may not have generated the data [4], [22].

Approaches to check if data invalidates the model have been developed for ordinary differential equation models of biological systems [21], [13], [23], [8], which are applicable when one wishes to model how the mean of a population of cells evolves over time. However, in many cases of interest, experiments are performed where a large population of cells is cultured, and at certain times the distribution of the species across the population is measured using techniques such as flow cytometry [2] or single cell RNA sequencing [12]. Such distributional data has been used for identifying the parameters of chemical reaction networks [18], [15], [17].

This work was supported in part by the U.S. AFOSR under grants FA9550-14-1-0060 and FA9550-22-1-0316.

¹T. W. Grunberg is with the Department of Electrical Engineering and Computer Science, MIT, Cambridge, MA 02139 USA grunberg@mit.edu

²D. Del Vecchio, is with the Department of Mechanical Engineering, MIT, Cambridge, MA 02139 USA ddv@mit.edu

However, the problem of model invalidation from distributional data is largely unexplored. A special case of distributional data is when only samples from the stationary distribution are measured, which is the case we focus on in this work. To allow for the possibility that the observed empirical distribution is perturbed in some way from the stationary distribution of the idealized chemical reaction network model, e.g. by outliers or finite sample effects, we study a *robust* model invalidation problem, that requires the observed distribution to be sufficiently far in total variation distance from all stationary distributions achievable by the model before we invalidate the model. In this setting, the model invalidation problem can be naturally formulated as a nonconvex quadratic program, where the nonconvexity arises through the bilinear interactions between the (unknown) parameters and the probability mass function through the chemical master equation. Such a quadratic program cannot be solved in practice, since the number of decision variables is equal to the number of relevant microstates, which even for small systems may be extremely large, or even infinite. Our approach summarizes the stationary distribution using finitely many “exponentially weighted moments,” which we introduce. In this way we obtain a much smaller, though still nonconvex, quadratic program that can be solved for biologically relevant systems using commercial software such as Gurobi [1].

Our approach uses methods that are similar to the semidefinite and linear programming approaches to bounding the moments of the stationary distribution [10], [7]. Such methods are based on the fact that the lower order moments must satisfy certain constraints, which come from the CME [9], and the fact that one can describe the set of moments achieved by any probability distribution using semidefinite programming [14]. However, when the maximum molecular counts of the species are unbounded, the total variation distance between two distributions cannot be bounded from below by comparing finitely many moments of the two distributions. This is due to the fact that small perturbations in total variation can cause unbounded changes in the moments, and results in an inability to robustly invalidate a model using finitely many moments. To resolve this issue, we introduce a novel family of generalized moments which are better suited to the model invalidation problem for chemical reaction networks in the sense that the total variation distance between two distributions can be underapproximated by comparing finitely many of our generalized moments.

The remainder of the paper is organized as follows. In Section II, we provide mathematical background and

introduce our definition of robust model invalidation. In Section III, we introduce these generalized moments and show how to construct sets where the first k generalized moments of the stationary distribution must lie. Then, in Section IV we exploit this construction to give a solution to the robust model invalidation problem. Section V we show a connection between robust model invalidation and hypothesis testing. We illustrate our method by applying it to an example chemical reaction network in Section VI, and provide concluding remarks and directions for future work in Section VII.

II. PROBLEM SETTING

A. The Chemical Master Equation

We consider a chemical reaction network (CRN) consisting of n species $\mathbf{X} = [X_1, \dots, X_n]^T$ interacting via r reactions $\langle \xi_r^i, \mathbf{X} \rangle \xrightarrow{\theta_i} \langle \xi_p^i, \mathbf{X} \rangle$ in a container with unit volume. When reaction i fires, the species counts \mathbf{X} changes by $\xi^i = \xi_p^i - \xi_r^i$. Each reaction occurs at rate $q_i(\mathbf{X})$, which is proportional to $\theta_i \geq 0$, the reaction rate constant of reaction i . Specifically, we consider CRNs with only zeroth, first, and second order reactions. A zeroth order reaction is of the form $\emptyset \xrightarrow{\theta_i} \langle \xi_p^i, \mathbf{X} \rangle$, and has propensity $q_i(\mathbf{X}) = \theta_i$. A first order reaction is of the form $X_s \xrightarrow{\theta_i} \langle \xi_p^i, \mathbf{X} \rangle$ for some s , and has propensity $q_i(\mathbf{X}) = \theta_i X_s$. A second order reaction is of the form $X_s + X_{s'} \xrightarrow{\theta_i} \langle \xi_p^i, \mathbf{X} \rangle$ for some s, s' , and has propensity $q_i(\mathbf{X}) = \theta_i X_s X_{s'}$ if $s \neq s'$, and propensity $q_i(\mathbf{X}) = \frac{1}{2} \theta_i X_s (X_s - 1)$ if $s = s'$. Since we assume that there are no third order or higher order reactions, for all i , $q_i(\mathbf{X})$ is a polynomial of degree less than or equal to two.

The Chemical Master Equation (CME) describes the time evolution of the probability mass function of the species counts:

$$\begin{aligned} \frac{d}{dt} \mathbb{P}[\mathbf{X}(t) = \mathbf{x}] &= \sum_{i=1}^r q_i(\mathbf{x} - \xi^i) \mathbb{P}[\mathbf{X}(t) = \mathbf{x} - \xi^i] \\ &\quad - \sum_{i=1}^r q_i(\mathbf{x}) \mathbb{P}[\mathbf{X}(t) = \mathbf{x}]. \end{aligned}$$

In this work we assume that $\mathbf{X}(t)$ is ergodic, and as such the CME has a unique steady state solution, which is the stationary distribution of $\mathbf{X}(t)$.

For any function $h: \mathbb{Z}_{\geq 0}^n \rightarrow \mathbb{R}$ we have the corresponding *generalized moment* $\mathbb{E}[h(\mathbf{X})]$, when the expectation exists. We have that the derivative of $\mathbb{E}[h(\mathbf{X})]$ is given by

$$\frac{d}{dt} \mathbb{E}[h(\mathbf{X})] = \mathbb{E} \left[\sum_{i=1}^r q_i(\mathbf{X}) (h(\mathbf{X} + \xi^i) - h(\mathbf{X})) \right]. \quad (1)$$

For a proof of this formula see Lemma 2.1 in [9].

B. Total variation distance

The *total variation distance* between two probability distributions $\nu(\mathbf{x})$ and $\rho(\mathbf{x})$ in $\mathcal{P}(\mathbb{Z}_{\geq 0}^n)$, the set of probability distributions over $\mathbb{Z}_{\geq 0}^n$, is defined as

$$d_{TV}(\nu, \rho) = \sup_{S \subseteq \mathbb{Z}_{\geq 0}^n} |\nu(S) - \rho(S)|$$

where $\nu(S) = \sum_{\mathbf{x} \in S} \nu(\mathbf{x})$ and $\rho(S) = \sum_{\mathbf{x} \in S} \rho(\mathbf{x})$.

C. Robust model invalidation

Given a CRN with unknown reaction rate constants $\theta \in \Theta$, let $\pi(\mathbf{x}; \theta)$ denote the stationary distribution. Given an observed distribution, the model is invalidated if no value of $\theta \in \Theta$ reproduces the observed distribution, i.e. if there does not exist a $\theta \in \Theta$ such that $\pi(\mathbf{x}; \theta)$ matches the observed distribution. To account for the possibility of measurement error, in this work we consider a robust variant of model invalidation, which we formally define as follows:

Definition 2.1: Given a CRN model with parameters $\theta \in \Theta$ and stationary distribution $\pi(\mathbf{x}; \theta)$, an observed distribution $\hat{\pi}(\mathbf{x})$, and a threshold $\eta \geq 0$, $\hat{\pi}(\mathbf{x})$ *robustly invalidates the model* $\pi(\mathbf{x}; \theta)$ with threshold η if there does not exist $\theta \in \Theta$ such that $d_{TV}(\pi(\cdot; \theta), \hat{\pi}) \leq \eta$.

Our notion of robust model invalidation can capture multiple effects. For example, if ζ fraction of samples are corrupted, e.g. drawn from an arbitrary distribution ν , the observed distribution with infinitely many samples is

$$\hat{\pi} = (1 - \zeta)\pi(\cdot; \theta) + \zeta\nu. \quad (2)$$

Since it follows that $d_{TV}(\hat{\pi}, \pi(\cdot; \theta)) \leq \zeta$, one can pick $\eta = \zeta$ and be guaranteed that if $\hat{\pi}$ robustly invalidates the model $\pi(\cdot; \theta)$ with threshold η , then $\hat{\pi}$ could not have been generated by (2) for any $\theta \in \Theta$ and any distribution ν . Alternatively, robust model invalidation can capture $\hat{\pi}$ being an arbitrary ‘‘perturbation’’ in total variation distance to $\pi(\cdot; \theta)$, representing the fact that CRN models are rarely perfect descriptions of biological systems. In Section V, we will see another motivation for our definition of robust model invalidation through its connection to hypothesis testing. The problem we consider in this work is constructing an algorithm for robust model invalidation, that is, given a CRN model $\pi(\cdot; \theta)$ with $\theta \in \Theta$, observed distribution $\hat{\pi}$, and a threshold η , determine if $\hat{\pi}$ robustly invalidates $\pi(\cdot; \theta)$. As shown in Figure 1, we seek an algorithm having the property that the model is never incorrectly robustly invalidated. We do however allow the algorithm to be conservative, i.e. it may fail to declare the model robustly invalidated even when $\hat{\pi}$ robustly invalidates the model according to Definition 2.1. This mirrors the philosophy behind model invalidation, where declaring the model robustly invalidated represents a firm belief that the model is incorrect, and the alternative is that we do not know if the model is correct or not. When using an algorithm such as the one we develop in this work, this lack of knowledge may be because $\hat{\pi}$ is close to $\pi(\cdot; \Theta)$, or it may be because $\hat{\pi}$ is not close to $\pi(\cdot; \Theta)$, but our algorithm could not prove this fact. As we will see, allowing for this type of conservatism allows us to construct a computationally efficient algorithm for robust model invalidation.

III. EXPONENTIALLY WEIGHTED MOMENTS

In this section, we introduce the generalized moments that are necessary to solve the model invalidation problem in Section IV. In Section III-A, we then show how to construct




Fig. 1: Model invalidation setup. \mathcal{P} is the space of all probability distributions. The set $\pi(\mathbf{x}; \Theta) = \{\pi(\mathbf{x}; \boldsymbol{\theta}) | \boldsymbol{\theta} \in \Theta\}$ is all distributions that correspond to some $\boldsymbol{\theta} \in \Theta$. The dash line shows the set of all distributions ν such that $d_{TV}(\pi(\cdot; \boldsymbol{\theta}), \nu) \leq \eta$ for some $\boldsymbol{\theta} \in \Theta$. Distribution $\hat{\pi}_1$ is in $\pi(\mathbf{x}; \Theta)$ and should not invalidate the model, $\hat{\pi}_2$ is not within η in total variation of $\pi(\mathbf{x}; \Theta)$ and so would ideally invalidate the model. Distribution $\hat{\pi}_3$ is within η of $\pi(\mathbf{x}; \Theta)$ in total variation, and so cannot invalidate the model. Distribution $\hat{\pi}_4$ is farther than η in total variation from $\pi(\mathbf{x}; \Theta)$ and so robustly invalidates the model. We seek an algorithm which can certify invalidity of the model, but we allow the algorithm to be conservative. This corresponds to finding the grey set \mathcal{I} , and an associated algorithm which can determine if the observed distribution is in \mathcal{I} or not. In the shown cartoon, only $\hat{\pi}_4$ would result in the algorithm robustly invalidating the model.

constraints that hold for the generalized moments of any distribution that is close to $\hat{\pi}$, and in Section III-B we use the CME to construct bilinear constraints relating the generalized moments of the stationary distribution and the reaction rate constants.

Let ν be a probability mass function over $\mathbb{Z}_{\geq 0}^n$ and let $\gamma > 0$ such that $\frac{1}{\gamma} \in \mathbb{Z}$. For $\mathbf{i} = (i_1, i_2, \dots, i_n) \in \mathbb{Z}_{\geq 0}^n$ we define

$$g_{\mathbf{i}}(\mathbf{x}) = c_{\mathbf{i}} e^{-\gamma \langle \mathbf{1}, \mathbf{x} \rangle} \prod_{j=1}^n x_j^{i_j},$$

where $c_{\mathbf{i}} = \left(e^{-\sum_{j=1}^n i_j} \prod_{j=1}^n \left(\frac{i_j}{\gamma} \right)^{i_j} \right)^{-1}$. Observe that $0 \leq g_{\mathbf{i}}(\mathbf{x}) \leq 1$ for all $\mathbf{x} \geq 0$, and that $g_{\mathbf{i}}(\mathbf{x})$ takes value 1 at exactly $\mathbf{x} = \frac{1}{\gamma} \mathbf{i}$. We then define the exponentially weighted moments of ν as

$$\lambda_{\mathbf{i}}(\nu) = \mathbb{E}_{\nu} [g_{\mathbf{i}}(\mathbf{X})].$$

Let $k \in \mathbb{N}$. For convenience let $\phi : \mathbb{Z}_{> 0} \rightarrow \mathbb{Z}_{\geq 0}^n$ be a one-to-one function such that $\phi(\{1, \dots, K\}) = \{0, \mathbf{1}, \dots, k\}^n$. We define

$$\mathbf{g}^k(\mathbf{x}) = [g_{\phi(1)}(\mathbf{x}) \quad g_{\phi(2)}(\mathbf{x}) \quad \dots \quad g_{\phi(K)}(\mathbf{x})]^T,$$

where $K = (k+1)^n$, and

$$\boldsymbol{\lambda}^k(\nu) = \mathbb{E}_{\nu} [\mathbf{g}^k(\mathbf{X})].$$

Thus, for a distribution ν , $\boldsymbol{\lambda}^k(\nu)$ is the vector of all generalized moments such that the degree in each species is $\leq k$, i.e. $\lambda_{\mathbf{i}}$ where $\mathbf{i} \leq k$. It will sometimes be convenient to express $\boldsymbol{\lambda}^k(\nu)$ using matrix notation. Let $\mathbf{p} \in \ell_1$ be a probability vector representing $\nu \in \mathcal{P}(\mathbb{Z}_{\geq 0}^n)$. We denote by \mathcal{G} the matrix with $(k+1)^n$ rows and infinite columns such that $\boldsymbol{\lambda}^k(\nu) = \mathcal{G}\mathbf{p}$. Observe that each column of \mathcal{G} will be given by $\mathbf{g}^k(\mathbf{x}^*)$ for some $\mathbf{x}^* \in \mathbb{Z}_{\geq 0}^n$.

A. Exponentially weighted moments of distributions close to $\hat{\pi}$

We now consider the set of values of $\boldsymbol{\lambda}^k(\nu)$ for some distribution ν that is within η of $\hat{\pi}$ in total variation. These values are given by the set $\Lambda^k(\hat{\pi}, \eta)$, defined formally as

$$\Lambda^k(\hat{\pi}, \eta) = \{ \boldsymbol{\lambda}^k(\nu) | \nu \in \mathcal{P}(\mathbb{Z}_{\geq 0}^n) \text{ s.t. } d_{TV}(\nu, \hat{\pi}) \leq \eta \}.$$

In this section we give one approach to producing an outer approximation of $\Lambda^k(\hat{\pi}, \eta)$. Let us define c_{λ} as the maximum 2-norm discrepancy between $\boldsymbol{\lambda}^k(\hat{\pi})$ and any vector of generalized moments of a distribution that is within η of $\hat{\pi}$ in total variation distance:

$$c_{\lambda}(\hat{\pi}, \eta) = \sup_{\nu \in \mathcal{P}(\mathbb{Z}_{\geq 0}^n) : d_{TV}(\nu, \hat{\pi}) \leq \eta} \| \boldsymbol{\lambda}^k(\nu) - \boldsymbol{\lambda}^k(\hat{\pi}) \|_2.$$

We have that $\{ \boldsymbol{\lambda} | \| \boldsymbol{\lambda}^k(\nu) - \boldsymbol{\lambda}^k(\hat{\pi}) \|_2 \leq c_{\lambda}(\hat{\pi}, \eta) \} \supseteq \Lambda^k(\hat{\pi}, \eta)$. We now show how to compute upper bounds on c_{λ} , which can then be used to create an outer approximation to $\Lambda^k(\hat{\pi}, \eta)$. We start by observing that

$$c_{\lambda}(\hat{\pi}, \eta) = \sup_{\mathbf{p} \in \ell_1 : \mathbf{p} \geq 0, \sum_i p_i = 1, \| \mathbf{p} - \hat{\mathbf{p}} \|_1 \leq 2\eta} \| \mathcal{G}(\mathbf{p} - \hat{\mathbf{p}}) \|_2, \quad (3)$$

where $\hat{\mathbf{p}}$ is the probability vector representation of $\hat{\pi}$. Define e_{λ} as

$$e_{\lambda} = \sup_{\tilde{\mathbf{p}} \in \ell_1 : \| \tilde{\mathbf{p}} \|_1 \leq 2\eta, \sum_i \tilde{p}_i = 0} \| \mathcal{G}\tilde{\mathbf{p}} \|_2. \quad (4)$$

Since the feasible set in (4) is a superset of the feasible set in (3), e_{λ} provides an upper bound on c_{λ} . Since (4) is a convex maximization problem, to compute e_{λ} we must check the value of $\| \mathcal{G}\tilde{\mathbf{p}} \|_2$ at every extreme point of the feasible set. This is formalized, with the extreme points given explicitly, in Lemma 3.1.

Lemma 3.1: Let $H = \{ \mathbf{p} \in \ell_1 | \sum_i p_i = 0, \| \mathbf{p} \|_1 \leq 2\eta \}$, $K \in \mathbb{R}^{m \times \infty}$ have finite absolute row sums, and $\mathbf{b} \in \mathbb{R}^m$. For $i, j \geq 1$, let $\mathbf{v}_{ij} = \eta \boldsymbol{\delta}_i - \eta \boldsymbol{\delta}_j$ where $\boldsymbol{\delta}_i$ is the i^{th} canonical unit vector in ℓ_1 . We have that

$$\sup_{\mathbf{p} \in H} \| K\mathbf{p} + \mathbf{b} \|_2 = \sup_{1 \leq i, j} \| K\mathbf{v}_{ij} + \mathbf{b} \|_2.$$

Proof: See appendix A. ■

Lemma 3.1 implies that

$$e_{\lambda} = \sup_{i, j \geq 1} \| \mathcal{G}\mathbf{v}_{ij} \|_2 = \sup_{\mathbf{x}, \mathbf{x}' \in \mathbb{Z}_{\geq 0}^n} \eta \| \mathbf{g}^k(\mathbf{x}) - \mathbf{g}^k(\mathbf{x}') \|_2,$$

where $\mathbf{v}_{ij} = \eta \boldsymbol{\delta}_i - \eta \boldsymbol{\delta}_j$. Partition $\mathbb{Z}_{\geq 0}^n$ into three sets,

$$\begin{aligned} S_{N'}^1 &= \{ (\mathbf{x}, \mathbf{x}') \in \mathbb{Z}_{\geq 0}^{2n} | \mathbf{x}, \mathbf{x}' \leq N' \}, \\ S_{N'}^2 &= \{ (\mathbf{x}, \mathbf{x}') \in \mathbb{Z}_{\geq 0}^{2n} | \mathbf{x}, \mathbf{x}' \not\leq N' \}, \\ S_{N'}^3 &= \{ (\mathbf{x}, \mathbf{x}') \in \mathbb{Z}_{\geq 0}^{2n} | \mathbf{x} \not\leq N' \text{ or } \mathbf{x}' \not\leq N', \text{ but not both} \}. \end{aligned}$$

Here, $\mathbf{x} \leq N'$ means that every elements of \mathbf{x} is less than or equal to N' , and $\mathbf{x} \not\leq N'$ means that at least one element of \mathbf{x} is greater than N' . Additionally, let

$$B_{N'} = \{\mathbf{x} \in \mathbb{Z}_{\geq 0}^n \mid \mathbf{x} \leq N' \text{ and } \exists i \text{ s.t. } x_i = N'\},$$

which the set of points in $\mathbb{Z}_{\geq 0}^n$ on the ‘‘boundary’’ of the subset where $\mathbf{x} \leq N'$. Let

$$e_\lambda^1(N') = \eta \cdot \max_{(\mathbf{x}, \mathbf{x}') \in S_{N'}^1} \|\mathbf{g}^k(\mathbf{x}) - \mathbf{g}^k(\mathbf{x}')\|_2,$$

$$e_\lambda^2(N') = \eta \cdot \sup_{(\mathbf{x}, \mathbf{x}') \in S_{N'}^2} \|\mathbf{g}^k(\mathbf{x}) - \mathbf{g}^k(\mathbf{x}')\|_2,$$

$$e_\lambda^3(N') = \eta \cdot \sup_{(\mathbf{x}, \mathbf{x}') \in S_{N'}^3} \|\mathbf{g}^k(\mathbf{x}) - \mathbf{g}^k(\mathbf{x}')\|_2.$$

Since $S_{N'}^1 \cup S_{N'}^2 \cup S_{N'}^3 = \mathbb{Z}_{\geq 0}^{2n}$, we have that

$$e_\lambda(N') = \max \{e_\lambda^1(N'), e_\lambda^2(N'), e_\lambda^3(N')\}.$$

One can compute $e_\lambda^1(N')$ since it is a maximum over a finite number of points. However, it is not obvious how to compute $e_\lambda^2(N')$ and $e_\lambda^3(N')$. Let

$$\tilde{e}_\lambda^2(N') = 2\eta \max_{\mathbf{x} \in B_{N'}} \|\mathbf{g}^k(\mathbf{x})\|_2,$$

and

$$\tilde{e}_\lambda^3(N') = \eta \cdot \max_{\mathbf{x} \in \mathbb{Z}_{\geq 0}^n, \mathbf{x} \leq N'} \|\mathbf{g}^k(\mathbf{x})\|_2 + \eta \cdot \max_{\mathbf{x} \in B_{N'}} \|\mathbf{g}^k(\mathbf{x})\|_2.$$

Define

$$\tilde{e}_\lambda(N') = \max \{e_\lambda^1(N'), \tilde{e}_\lambda^2(N'), \tilde{e}_\lambda^3(N')\}.$$

For sufficiently large N' , $e_\lambda \leq \tilde{e}_\lambda(N')$, with the bound becoming tight as $N' \rightarrow \infty$. This is formalized in the following lemma.

Lemma 3.2: For any integer $N' > k/\gamma$ we have that

$$e_\lambda \leq \tilde{e}_\lambda(N').$$

Furthermore, $\lim_{N' \rightarrow \infty} \tilde{e}_\lambda(N') = e_\lambda$.

Proof: See Appendix B. \blacksquare

Based on Lemma 3.2 we can construct an outer approximation to $\Lambda^k(\hat{\pi}, \eta)$, as given in the following theorem.

Theorem 3.1: For any integer $N' > k/\gamma$ we have that

$$\{\boldsymbol{\lambda} \mid \|\boldsymbol{\lambda} - \boldsymbol{\lambda}^k(\hat{\pi})\|_2 \leq \tilde{e}_\lambda(N')\} \supseteq \Lambda^k(\hat{\pi}, \eta).$$

Proof: Observe that $c_\lambda \leq e_\lambda$. The result then follows from Lemma 3.2. \blacksquare

We now give another type of constraint that $\boldsymbol{\lambda}^k(\nu)$ satisfies for all ν such that $d_{TV}(\nu, \hat{\pi}) \leq \eta$. We exploit that fact that for certain values of γ and k , the elements of $\boldsymbol{\lambda}^k$ are constrained to be approximately in a lower dimensional subspace of $\mathbb{R}^{(k+1)^n}$. If we can find a $T \in \mathbb{R}^{m \times (k+1)^n}$ such that the nullspace of T is the subspace of $\mathbb{R}^{(k+1)^n}$ that $\boldsymbol{\lambda}^k$ is close to, $\|T\boldsymbol{\lambda}^k\|_2$ will be small. We first give a general method for bounding

$$c_T = \sup_{\nu \in \mathcal{P}(\mathbb{Z}_{\geq 0}^n): d_{TV}(\nu, \hat{\pi}) \leq \eta} \|T\boldsymbol{\lambda}^k(\nu)\|_2 \quad (5)$$

for an arbitrary T . Then, we give a heuristic for selecting T . Our development of the bound on c_T proceeds similarly to our bound on c_λ . Let

$$e_T = \sup_{\mathbf{p}: \sum_i p_i = 0, \|\mathbf{p} - \hat{\mathbf{p}}\|_1 \leq 2\eta} \|T\mathcal{G}\mathbf{p}\|_2. \quad (6)$$

From the fact that the feasible set is larger in (6) than in (5) we have that $c_T \leq e_T$. An equivalent expression for e_T is

$$e_T = \sup_{\hat{\mathbf{p}}: \sum_i \hat{p}_i = 0, \|\hat{\mathbf{p}}\|_1 \leq 2\eta} \|T\boldsymbol{\lambda}^k(\hat{\pi}) + T\mathcal{G}\hat{\mathbf{p}}\|_2.$$

By using Lemma 3.1 we can write e_T as

$$\begin{aligned} e_T &= \sup_{i, j \in \mathbb{Z}_{> 0}} \|T\boldsymbol{\lambda}^k(\hat{\pi}) + T\mathbf{v}_{ij}\|_2, \\ &= \sup_{\mathbf{x}, \mathbf{x}' \in \mathbb{Z}_{\geq 0}^n} \|T\boldsymbol{\lambda}^k(\hat{\pi}) + \eta T(\mathbf{g}^k(\mathbf{x}) - \mathbf{g}^k(\mathbf{x}'))\|_2. \end{aligned}$$

Similarly to our approach to bounding e_λ we define

$$e_T^1(N') = \max_{(\mathbf{x}, \mathbf{x}') \in S_{N'}^1} \|T\boldsymbol{\lambda}^k(\hat{\pi}) + \eta T(\mathbf{g}^k(\mathbf{x}) - \mathbf{g}^k(\mathbf{x}'))\|_2,$$

$$e_T^2(N') = \sup_{(\mathbf{x}, \mathbf{x}') \in S_{N'}^2} \|T\boldsymbol{\lambda}^k(\hat{\pi}) + \eta T(\mathbf{g}^k(\mathbf{x}) - \mathbf{g}^k(\mathbf{x}'))\|_2,$$

$$e_T^3(N') = \sup_{(\mathbf{x}, \mathbf{x}') \in S_{N'}^3} \|T\boldsymbol{\lambda}^k(\hat{\pi}) + \eta T(\mathbf{g}^k(\mathbf{x}) - \mathbf{g}^k(\mathbf{x}'))\|_2.$$

Since $S_{N'}^1 \cup S_{N'}^2 \cup S_{N'}^3 = \mathbb{Z}_{\geq 0}^{2n}$, we have that

$$e_T(N') = \max \{e_T^1(N'), e_T^2(N'), e_T^3(N')\}.$$

One can compute $e_T^1(N')$, but we must bound $e_T^2(N')$ and $e_T^3(N')$. Let

$$\tilde{e}_T^2(N') = \|T\boldsymbol{\lambda}^k(\hat{\pi})\|_2 + 2\eta \max_{\mathbf{x} \in B_{N'}} \|T\mathbf{g}^k(\mathbf{x})\|_2$$

and

$$\begin{aligned} \tilde{e}_T^3(N') &= \max_{\mathbf{x} \in \mathbb{Z}_{\geq 0}^n, \mathbf{x} \leq N', \kappa \in \{+1, -1\}} \|T\boldsymbol{\lambda}^k(\hat{\pi}) \\ &\quad + \kappa \eta T\mathbf{g}^k(\mathbf{x})\|_2 + \eta \cdot \max_{\mathbf{x} \in B_{N'}} \|T\mathbf{g}^k(\mathbf{x})\|_2. \end{aligned}$$

We define

$$\tilde{e}_T(N') = \max \{e_T^1(N'), \tilde{e}_T^2(N'), \tilde{e}_T^3(N')\},$$

and have the following lemma, which establishes that $\tilde{e}_T(N')$, with N' sufficiently large, can be used to bound e_T .

Lemma 3.3: For any integer $N' > k/\gamma$ we have that

$$e_T \leq \tilde{e}_T(N').$$

Furthermore, $\lim_{N' \rightarrow \infty} \tilde{e}_T(N') = e_T$.

Proof: See Appendix B. \blacksquare

Based on Lemma 3.3 we can construct an outer approximation of $\Lambda^k(\hat{\pi}, \eta)$, which we formalize in the following theorem.

Theorem 3.2: For any integer $N' > k/\gamma$ we have that

$$\{\boldsymbol{\lambda} \mid \|T\boldsymbol{\lambda}\|_2 \leq \tilde{e}_T(N')\} \supseteq \Lambda^k(\hat{\pi}, \eta).$$

Proof: Observe that $c_T \leq e_T$. The result then follows from Lemma 3.3. \blacksquare

While there are many methods to select a matrix T , and Theorem 3.2 will hold for any matrix T , we wish to choose a T so that the intersection of the sets given in Theorems 3.1 and 3.2 is a good outer approximation of $\Lambda^k(\hat{\pi}, \eta)$. Here we simply give a heuristic for selecting T , which is inspired by moment closure techniques for the standard moments [19]. Let λ^k be partitioned into $\lambda^k = [(\lambda^{0:k-2})^T \quad (\lambda^{k-1:k})^T]^T$, where $\lambda^{0:k-2}$ is composed of the elements corresponding to $g_i(\mathbf{x})$ with $i \leq k-2$ and $\lambda^{k-1:k}$ is composed of all the other elements. Note that the existence of such a partitioning requires a particular choice of ϕ . For some choices of γ and k , it will be true that there exists $\tilde{T} \in \mathbb{R}^{\dim \lambda^{k-1:k} \times (k+1)^n}$ such that

$$\forall \nu \in \mathcal{P}(\mathbb{Z}_{\geq 0}^n) \text{ s.t. } d_{TV}(\nu, \hat{\pi}) \leq \eta, \lambda^{k-1:k} \approx \tilde{T} \lambda^{0:k-2}.$$

This expresses the idea that at least for distributions close to $\hat{\pi}$, one can approximate $\lambda^{k-1:k}(\nu)$ by a linear combination of the elements of $\lambda^{0:k-2}$. A heuristic for picking such a \tilde{T} is based on least squares,

$$\tilde{T} = G_{N''}^{0:k-1} (G_{N''}^{k-1:k})^\dagger,$$

where $G_{N''} = \begin{bmatrix} G_{N''}^{0:k-1} \\ G_{N''}^{k-1:k} \end{bmatrix}$ is the matrix formed from the columns of \mathcal{G} corresponding to $\mathbf{x} \leq N''$, and N'' is picked such that most of the probability mass of $\hat{\pi}$ is contained in the region $\mathbf{x} \leq N''$. We then pick $T = [\tilde{T} \quad I]$. While we have not justified that the approximation will be good, the correctness of the robust model invalidation method we introduce in Section IV depends only upon Theorem 3.2, which holds for any matrix T , and thus using a heuristic for selecting T does not affect the rigor in our results.

Remark 3.1: An alternative approach to selecting T is to take the singular value decomposition of $G_{N''}$. If

$$G_{N''} = [U_1 \quad U_2] \Sigma V^T,$$

where the columns of U_2 correspond to the small singular values, we expect that $\|U_2^T \lambda^k(\nu)\|_2$ will be small for all ν , and thus $T = U_2^T$ is a reasonable choice of T . We defer further study of the selection of T for future work.

B. Exponentially weighted moments of the stationary distribution

Here we show that when $\pi(\mathbf{x}; \boldsymbol{\theta})$ is the stationary distribution of the CME, $\lambda^k(\pi(\cdot; \boldsymbol{\theta}))$ must satisfy particular constraints. Specifically, letting $\omega(\mathbf{x}; t, \boldsymbol{\theta})$ be the solution to the CME, we define $A^k(\boldsymbol{\theta}) = \sum_{i=1}^r \theta_i A_i^k$ as the matrix valued function such that

$$\frac{d}{dt} \lambda^{k-2}(\omega(\cdot; t, \boldsymbol{\theta})) = A^k(\boldsymbol{\theta}) \lambda^k(\omega(\cdot; t, \boldsymbol{\theta})), \quad (7)$$

when such an $A^k(\boldsymbol{\theta})$ exists. For the stationary distribution $\pi(\mathbf{x}; \boldsymbol{\theta})$, setting $0 = \frac{d}{dt} \lambda^{k-2}(\pi(\cdot; \boldsymbol{\theta}))$ we have the constraints

$$0 = A^k(\boldsymbol{\theta}) \lambda^k(\pi(\cdot; \boldsymbol{\theta})). \quad (8)$$

We stress that the existence of such constraints is a property of the particular generalized moments that we have chosen,

and in general, for a set of functions $g_i(\mathbf{x})$, no such constraints exist. However, a similar property does hold for the standard moments, i.e. averages with respect to monomials in \mathbf{X} [10], [7].

Theorem 3.3: Let $\omega(\mathbf{x}; t, \boldsymbol{\theta})$ be the solution to the CME, and let $\pi(\mathbf{x}; \boldsymbol{\theta})$ be the stationary distribution of the CME, both with reaction rate constants $\boldsymbol{\theta}$. Let $\lambda^k(\pi(\cdot; \boldsymbol{\theta}))$ be the vector of generalized moments corresponding to $\pi(\mathbf{x}; \boldsymbol{\theta})$. Then, there exists $A^k(\boldsymbol{\theta}) \neq 0$ linear in $\boldsymbol{\theta}$ satisfying (7) and (8) for all $\boldsymbol{\theta}$ such that $\mathbf{X}(t)$ is ergodic.

Proof: Let $\omega(\mathbf{x}; t, \boldsymbol{\theta})$ be the distribution of $\mathbf{X}(t)$. We show that $A^k(\boldsymbol{\theta})$ can be obtained by writing the derivatives of $\lambda^{k-2}(t)$ as a linear combination of the elements of $\lambda^k(t)$. For any multi index $\mathbf{i} \leq k-2$, we have from (1) that

$$\begin{aligned} \frac{d\lambda_{\mathbf{i}}(\omega(\cdot; t, \boldsymbol{\theta}))}{dt} &= \mathbb{E} \left[\sum_{j=1}^r q_j(\mathbf{X}) (g_{\mathbf{i}}(\mathbf{X} + \boldsymbol{\xi}^j) - g_{\mathbf{i}}(\mathbf{X})) \right], \\ &= \mathbb{E} \left[\sum_{j=1}^r q_j(\mathbf{X}) \left(c_i e^{-\gamma \langle \mathbf{1}, \mathbf{X} + \boldsymbol{\xi}^j \rangle} \prod_{s=1}^n (X_s + \xi_s^j)^{i_s} \right. \right. \\ &\quad \left. \left. - c_i e^{-\gamma \langle \mathbf{1}, \mathbf{X} \rangle} \prod_{s=1}^n (X_s)^{i_s} \right) \right], \\ &= \mathbb{E} \left[c_i e^{-\gamma \langle \mathbf{1}, \mathbf{X} \rangle} \sum_{j=1}^r q_j(\mathbf{X}) \left(e^{-\gamma \langle \mathbf{1}, \boldsymbol{\xi}^j \rangle} \prod_{s=1}^n (X_s + \xi_s^j)^{i_s} \right. \right. \\ &\quad \left. \left. - \prod_{s=1}^n (X_s)^{i_s} \right) \right], \end{aligned}$$

where all expectations are with respect to $\mathbf{X} \sim \omega(\cdot; t, \boldsymbol{\theta})$. One can see that under our assumption that $q_j(\mathbf{X})$ is a polynomial of order 2 or less, $\frac{d}{dt} \lambda_{\mathbf{i}}(t)$ can be written as a linear combination of the elements of λ^k . Additionally, since $q_j(\mathbf{X})$ is linear in θ_j , one can see that $\frac{d}{dt} \lambda_{\mathbf{i}}(t)$ is linear in $\boldsymbol{\theta}$. Thus, but setting $0 = \frac{d}{dt} \lambda_{\mathbf{i}}(t)$ for $0 \leq \mathbf{i} \leq k-2$, justified by the boundedness of $\mathbf{g}^k(\mathbf{x})$ and the ergodicity of $\mathbf{X}(t)$, we have the desired result for the generalized moments of the stationary distribution. \blacksquare

IV. ROBUST MODEL INVALIDATION

In this section we present the main result of this work, an algorithm that can certify that an observed distribution is farther than η from the model class $p(\mathbf{x}; \boldsymbol{\theta})$ in total variation. We consider optimization problem (10), which is a nonconvex quadratically constrained program due to the $0 = (\sum_{i=1}^r \theta_i A_i^k) \lambda$ constraint. The decision variables are $\boldsymbol{\theta} \in \mathbb{R}^r$ and $\lambda \in \mathbb{R}^{(k+1)^n}$, and we assume for simplicity that $\Theta \subseteq \mathbb{R}_{>0}^r$ is a polyhedral set.

$$\|\lambda - \lambda^k(\hat{\pi})\|_2 \leq \tilde{\epsilon}_\lambda(N') \quad (10a)$$

$$\|T\lambda\|_2 \leq \tilde{\epsilon}_T(N') \quad (10b)$$

$$0 = \left(\sum_{i=1}^r \theta_i A_i^k \right) \lambda \quad (10c)$$

$$\boldsymbol{\theta} \in \Theta \quad (10d)$$

As formalized in Theorem 4.1, if we can show infeasibility of (10), then we have a method to robustly invalidate a model. We note that despite the fact that (10) is nonconvex, it has only $r + (k+1)^n$ variables, and thus the size of (10) can be controlled by our choice of k . We will see in an example that $(k+1)^n$ can be made far smaller than the number of variables needed to naively represent $\pi(\mathbf{x}; \boldsymbol{\theta})$ to within ϵ in total variation distance. Nonconvex quadratically constrained feasibility problems can be solved by commercial solvers such as Gurobi, using a spatial branch and bound method [1].

Theorem 4.1: Consider a CRN with stationary distribution $\pi(\cdot; \boldsymbol{\theta})$ and parameters $\boldsymbol{\theta} \in \Theta$ and let $\eta \geq 0$ be a threshold. Let $N' > k/\gamma$. For an observed distribution $\hat{\pi}$, if (10) is infeasible, then the model $\pi(\mathbf{x}; \boldsymbol{\theta})$ is robustly invalidated with threshold η by $\hat{\pi}$.

Proof: We prove the contrapositive. Suppose the model $\pi(\mathbf{x}; \boldsymbol{\theta})$ is not robustly invalidated with threshold η by $\hat{\pi}$. Then, exists $\boldsymbol{\theta}^* \in \Theta$ such that $d_{TV}(\pi(\cdot; \boldsymbol{\theta}^*), \hat{\pi}) \leq \eta$. We show that (10) is feasible. We have that $\|\boldsymbol{\lambda}^k(\pi(\cdot; \boldsymbol{\theta}^*)) - \boldsymbol{\lambda}^k(\hat{\pi})\|_2 \leq \tilde{e}_\lambda(N')$ and $\|T\boldsymbol{\lambda}^k(\pi(\cdot; \boldsymbol{\theta}^*))\|_2 \leq \tilde{e}_T(N')$ from Theorems 3.1 and 3.2. Additionally, since $\pi(\cdot; \boldsymbol{\theta}^*)$ is the stationary distribution of the CRN with parameters $\boldsymbol{\theta}^*$, we have that $0 = (\sum_{i=1}^r \theta_i A_i^k) \boldsymbol{\lambda}^k(\pi(\cdot; \boldsymbol{\theta}^*))$ from Theorem 3.3. Hence, $\boldsymbol{\lambda}^k(\pi(\cdot; \boldsymbol{\theta}^*))$ is a solution to (10), which completes the proof. ■

Remark 4.1: A formulation of robust model invalidation similar to (10), but using the standard moments instead of the exponentially weighted moments is not possible. In fact, let $h_i(\mathbf{x}) = \prod_{j=1}^n x_j^{i_j}$, and define

$$\boldsymbol{\mu}^k(\nu) = \mathbb{E}_\nu \left[\begin{bmatrix} h_{\phi(1)}(\mathbf{X}) & h_{\phi(2)}(\mathbf{X}) & \dots & h_{\phi(K)}(\mathbf{X}) \end{bmatrix}^T \right],$$

the vector of standard moments of ν with $i \leq k$. For any $\eta > 0$,

$$\sup_{d_{TV}(\nu, \hat{\pi}) \leq \eta} \|\boldsymbol{\mu}^k(\nu) - \boldsymbol{\mu}^k(\hat{\pi})\|_2 = \infty,$$

which prevents us from constructing constraints analogous to (10a) using the standard moments.

V. RELATIONSHIP BETWEEN ROBUST MODEL INVALIDATION AND HYPOTHESIS TESTING

Our definition of robust model invalidation can also capture finite sample effects. To make this idea precise we consider a hypothesis testing setup. In this setup we measure N_s samples, $\mathbf{Y} = (\mathbf{y}^1, \mathbf{y}^2, \dots, \mathbf{y}^{N_s})$ i.i.d. from some distribution ν . We denote the joint distribution of \mathbf{Y} by ν^{N_s} . The null hypothesis H_0 is that the data was generated by our model, i.e. $\nu \in \pi(\cdot; \Theta)$. The alternative hypothesis H_1 is that $\nu \notin \pi(\cdot; \Theta)$. We denote the empirical distribution of the N_s samples by $\hat{p}(\mathbf{x}; \mathbf{Y}) = \frac{1}{N_s} \sum_{i=1}^{N_s} \mathbb{1}_{\mathbf{y}^i}(\mathbf{x})$. Suppose we are interested in constructing a decision rule that has type I error, rejecting H_0 when it is true, rate below a specified threshold, P_e . In this section, we show that any method of solving the robust model invalidation problem can, by an appropriate choice of $\eta = \eta(P_e, N_s)$, be used as a decision rule for the hypothesis testing problem with type I error rate

less than P_e . To make this precise, we define a decision rule as a function $R : (\mathbb{Z}_{\geq 0}^n)^{N_s} \rightarrow \{H_0, H_1\}$, so that

$$P_I(\boldsymbol{\theta}) = \mathbb{P}_{\mathbf{Y} \sim \pi^{N_s}(\cdot; \boldsymbol{\theta})} [R(\mathbf{Y}) = H_1]$$

is the type I error rate for decision rule R when the true model is $\pi(\cdot; \boldsymbol{\theta})$. We are interested in the worst case type I error rate, $P_I^* = \sup_{\boldsymbol{\theta} \in \Theta} P_I(\boldsymbol{\theta})$.

We say that a decision rule R is a *reliable rule for robust model invalidation with threshold η* if it satisfies

$$R^{-1}(H_1) \subseteq \left\{ \mathbf{Y} \mid \inf_{\boldsymbol{\theta} \in \Theta} d_{TV}(\hat{p}(\cdot; \mathbf{Y}), \pi(\cdot; \boldsymbol{\theta})) > \eta \right\}.$$

Remark 5.1: Any algorithm for robust model invalidation can be considered a decision rule, specifically,

$$R(\mathbf{Y}) = \begin{cases} H_1, & \hat{p}(\cdot; \mathbf{Y}) \text{ results in "invalidated"} \\ H_0, & \text{otherwise.} \end{cases}$$

Any decision rule derived in this way is a reliable rule for robust model invalidation with threshold η .

We now give a result that allows one to start with a algorithm for robust model invalidation parameterized by η , and for any desired type I error rate P_e , and number of samples N_s , create a decision rule that satisfies $P_I^* \leq P_e$ by setting the threshold η .

Theorem 5.1: Consider a model $\pi(\cdot; \boldsymbol{\theta})$ with $\boldsymbol{\theta} \in \Theta$, and let $\beta(N')$ be a strict monotonically decreasing function such that that for all $\boldsymbol{\theta} \in \Theta$, $\beta(N')$ satisfies

$$\sum_{\mathbf{x} \notin N'} \pi(\mathbf{x}; \boldsymbol{\theta}) \leq \beta(N').$$

Let $0 < P_e < 1$. Let $\eta' > 0$ be such that

$$\left(2^{(\lceil \beta^{-1}(\eta'/3) \rceil + 1)^n} - 2 \right) e^{-\frac{2}{9} N_s \eta'^2} \leq P_e.$$

Any decision rule R that is a reliable rule for robust model invalidation with threshold η' has type I error rate less than or equal to P_e , i.e. $P_I^* \leq P_e$.

Remark 5.2: In light of Remark 5.1 and Theorem 5.1, the model invalidation method developed in Section IV can be used to create a decision rule for the hypothesis testing framework with a specified upper bound on the type I error rate by selecting the threshold η as a function of the number of samples N_s . However, one must produce the function $\beta(N')$.

Proof: Let R be a decision rule that is a reliable rule for robust model invalidation with threshold η' . Then, for an arbitrary $\boldsymbol{\theta} \in \Theta$ we have

$$P_I(\boldsymbol{\theta}) \leq \mathbb{P}_{\mathbf{Y} \sim \pi^{N_s}(\cdot; \boldsymbol{\theta})} \left[\inf_{\boldsymbol{\theta}' \in \Theta} d_{TV}(\hat{p}(\cdot; \mathbf{Y}), \pi(\cdot; \boldsymbol{\theta}')) > \eta' \right], \quad (11a)$$

$$\leq \mathbb{P}_{\mathbf{Y} \sim \pi^{N_s}(\cdot; \boldsymbol{\theta})} [d_{TV}(\hat{p}(\cdot; \mathbf{Y}), \pi(\cdot; \boldsymbol{\theta})) > \eta']. \quad (11b)$$

To obtain an upper bound on the right hand side of (11b) we use large deviations theory. Specifically, from [11, Lemma 2], we have that for an arbitrary distribution ν , if \mathcal{B} is a finite set such that

$$\sum_{\mathbf{x} \notin \mathcal{B}} \nu(\mathbf{x}) \leq \frac{\eta'}{3},$$

then

$$\mathbb{P}_{\mathbf{Y} \sim \nu^{N_s}} [d_{TV}(\hat{p}(\cdot; \mathbf{Y}), \nu) > \eta'] \leq (2^{|\mathcal{B}|} - 2) e^{-\frac{2}{9} N_s \eta'^2}.$$

Observe that for all $\boldsymbol{\theta} \in \Theta$ the set $\mathcal{B} = \{\mathbf{x} \in \mathbb{Z}_{\geq 0}^n \mid \mathbf{x} \leq \lceil \beta^{-1}(\eta'/3) \rceil\}$ satisfies

$$\sum_{\mathbf{x} \notin \mathcal{B}} \pi(\mathbf{x}; \boldsymbol{\theta}) \leq \frac{\eta'}{3},$$

and that $|\mathcal{B}| = (\lceil \beta^{-1}(\eta'/3) \rceil + 1)^n$. Therefore, choosing $\nu = \pi(\cdot; \boldsymbol{\theta})$ for any $\boldsymbol{\theta} \in \Theta$ yields


$$\begin{aligned} \mathbb{P}_{\mathbf{Y} \sim \pi^{N_s}(\cdot; \boldsymbol{\theta})} [d_{TV}(\hat{p}(\cdot; \mathbf{Y}), \pi(\cdot; \boldsymbol{\theta})) > \eta'] \\ \leq \left(2^{(\lceil \beta^{-1}(\eta'/3) \rceil + 1)^n} - 2 \right) e^{-\frac{2}{9} N_s \eta'^2}, \end{aligned}$$

and by assumption the right hand side is bounded from above by P_e . Therefore, from (11b) we have that $P_I(\boldsymbol{\theta}) \leq P_e$, and hence $P_I^* = \sup_{\boldsymbol{\theta} \in \Theta} P_I(\boldsymbol{\theta}) \leq P_e$. ■

Remark 5.3: The approach given in this section uses the robust model invalidation method presented in Section IV to construct a decision rule. In this decision rule we select the threshold η a priori, i.e. it does not depend on \mathbf{Y} . An approach that might yield decision rules with the same type I error rate but a lower type II error (selecting H_0 erroneously) rate might be obtained by selecting the threshold as a function of \mathbf{Y} . We defer investigating such a possibility to future work.

VI. EXAMPLE

Here we present an example of our model invalidation framework applied to two different CRNs with two species, where our goal is to determine which of the two CRNs are consistent with observed data, and thus determine if the two molecules in the system are suitable for constructing a particular biomolecular circuit. One way to implement an integral controller using chemical reactions is with the *antithetic motif*, where the two controller species annihilate one another [5], [20]. In order to obtain integral control using the antithetic motif, it is critical that the two species truly annihilate one another, instead of interacting via a different mechanism. We consider having measured data from the stationary distribution of a system with two species, X_1 and X_2 , where it is possible that X_1 and X_2 interact through mutual degradation or that they interact by X_2 enzymatically degrading X_1 . Let \mathcal{R}_1 be



with stationary distribution $\pi_1(\cdot; \boldsymbol{\theta})$ and associated set for the parameters of $\Theta_1 = \{\boldsymbol{\theta} \in \mathbb{R}^5 \mid 1 \leq \boldsymbol{\theta} \leq 500\}$. The CRN \mathcal{R}_1 models the situation where X_1 and X_2 undergo mutual





Fig. 2: Measured distribution $\hat{\pi}$ for the example. The distribution $\hat{\pi}$, with the probability mass function shown as a heatmap in (a) and an isometric plot in (b). The distribution is $\hat{\pi} = 0.995\pi_1(\cdot; \boldsymbol{\theta}) + 0.005\mathbb{1}_{x_1=9, x_2=49}$, a mixture of the stationary distribution of \mathcal{R}_1 and the distribution $\mathbb{1}_{x_1=9, x_2=49}(\mathbf{x})$. The parameters for \mathcal{R}_1 are taken as $\boldsymbol{\theta} = [200, 1, 200, 1, 1]$.

degradation. Let \mathcal{R}_2 be



with stationary distribution $\pi_2(\cdot; \boldsymbol{\theta})$ and associated set for the parameters of $\Theta_2 = \{\boldsymbol{\theta} \in \mathbb{R}^5 \mid 1 \leq \boldsymbol{\theta} \leq 500\}$. The CRN \mathcal{R}_2 models the situation where X_2 enzymatically degrades X_1 . Consider the observed distribution $\hat{\pi}$ shown in Figure 2, generated by \mathcal{R}_1 mixed with outliers at a point. We select a threshold $\eta = 0.005$, and choose $k = 9$ and $\gamma = k/30$. We choose T according to the moment closure heuristic computed with $N'' = 50$. To solve the optimization problems for this example we use the MATLAB toolbox YALMIP [16] to set up the optimization problems, and Gurobi [1] to solve them. Constructing feasibility problem (10) for \mathcal{R}_1 , we find that (10) is feasible, and thus $\hat{\pi}$ does not robustly invalidate the model $\pi_1(\cdot; \boldsymbol{\theta})$. On the other hand, for \mathcal{R}_2 , (10) is infeasible, and thus by Theorem 4.1 $\hat{\pi}$ robustly invalidates model $\pi_2(\cdot; \boldsymbol{\theta})$ with threshold $\eta = 0.005$. We conclude that \mathcal{R}_2 is not the correct model for the interaction of X_1 and X_2 , and that \mathcal{R}_1 is a possible model for the interaction. If we know that the true model is either \mathcal{R}_1 or η up to a perturbation of size η in total variation distance, then we can conclude that X_1 and X_2 undergo mutual degradation and thus are suitable for constructing an integral controller.

VII. CONCLUSION

In this work, we studied the problem of model invalidation for CRNs from the stationary distribution. Our approach makes use of exponentially weighted moments, which we introduce, to certify that the observed distribution cannot be reproduced to within the specified tolerance in total variation by any value of the parameters. In this way we obtain a relatively small, though nonconvex, quadratic feasibility problem that can be solved by commercial software [1]. We

related robust model invalidation to hypothesis testing by showing that our model invalidation method can be used to construct decision rules with a specified type I error. As we have seen in the example, our method is practical for determining which CRN models are consistent with observed data. Future work includes reducing the conservatism of our proposed method by considering alternative methods to construct an outer approximation of $\Lambda^k(\hat{\pi}, \eta)$, and extensions to the case where only certain species in the system are measured.

REFERENCES

- [1] Gurobi optimizer reference manual version 9.5. <https://www.gurobi.com>, 2022.
- [2] Aysun Adan, Günel Alizada, Yağmur Kiraz, Yusuf Baran, and Ayten Nalbant. Flow cytometry: basic principles and applications. *Critical reviews in biotechnology*, 37(2):163–176, 2017.
- [3] David F Anderson and Thomas G Kurtz. Continuous time markov chain models for chemical reaction networks. In *Design and analysis of biomolecular circuits*, pages 3–42. Springer, 2011.
- [4] James Anderson and Antonis Papachristodoulou. On validation and invalidation of biological models. *BMC bioinformatics*, 10(1):1–13, 2009.
- [5] Stephanie K. Aoki, Gabriele Lillacci, Ankit Gupta, Armin Baumschlager, David Schweingruber, and Mustafa Khammash. A universal biomolecular integral feedback controller for robust perfect adaptation. *Nature*, 570(7762):533–537, 2019.
- [6] Felipe A Campos, Simone Bruno, Yi Fu, Domitilla Del Vecchio, and Ruth J Williams. Comparison theorems for stochastic chemical reaction networks. *arXiv preprint arXiv:2302.03091*, 2023.
- [7] 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.
- [8] Hana El-Samad, Stephen Prajna, Antonis Papachristodoulou, Mustafa Khammash, and John C Doyle. Model validation and robust stability analysis of the bacterial heat shock response using sostoools. In *42nd IEEE International Conference on Decision and Control (IEEE Cat. No. 03CH37475)*, volume 4, pages 3766–3771. IEEE, 2003.
- [9] Stefan Engblom. Computing the moments of high dimensional solutions of the master equation. *Applied Mathematics and Computation*, 180(2):498–515, 2006.
- [10] 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.
- [11] Siu-Wai Ho and Raymond W Yeung. On information divergence measures and a unified typicality. *IEEE Transactions on Information Theory*, 56(12):5893–5905, 2010.
- [12] Aleksandra A Kolodziejczyk, Jong Kyoung Kim, Valentine Svensson, John C Marioni, and Sarah A Teichmann. The technology and biology of single-cell rna sequencing. *Molecular cell*, 58(4):610–620, 2015.
- [13] Lars Kuepfer, Uwe Sauer, and Pablo A. Parrilo. Efficient classification of complete parameter regions based on semidefinite programming. *BMC Bioinformatics*, 8(1):12, 2007.
- [14] Jean Bernard Lasserre. *Moments, positive polynomials and their applications*, volume 1. World Scientific, 2009.
- [15] Gabriele Lillacci and Mustafa Khammash. The signal within the noise: efficient inference of stochastic gene regulation models using fluorescence histograms and stochastic simulations. *Bioinformatics*, 29(18):2311–2319, 6/26/2021 2013.
- [16] J. Löfberg. Yalmip : A toolbox for modeling and optimization in matlab. In *In Proceedings of the CACSD Conference*, Taipei, Taiwan, 2004.
- [17] Brian Munsy, Guoliang Li, Zachary R. Fox, Douglas P. Shepherd, and Gregor Neuert. Distribution shapes govern the discovery of predictive models for gene regulation. *Proceedings of the National Academy of Sciences*, 115(29):7533, 07 2018.
- [18] Brian Munsy, Brooke Trinh, and Mustafa Khammash. Listening to the noise: random fluctuations reveal gene network parameters. *Molecular Systems Biology*, 5(1):318, 2020/09/10 2009.

- [19] Mohammad Naghnaeian and Domitilla Del Vecchio. Robust moment closure method for the chemical master equation. In *2017 IEEE Conference on Control Technology and Applications (CCTA)*, pages 967–972. IEEE, 2017.
- [20] Yili Qian and Domitilla Del Vecchio. Realizing ‘integral control’ in living cells: how to overcome leaky integration due to dilution? *Journal of The Royal Society Interface*, 15(139), 2018.
- [21] Philipp Rumschinski, Steffen Borchers, Sandro Bosio, Robert Weismantel, and Rolf Findeisen. Set-base dynamical parameter estimation and model invalidation for biochemical reaction networks. *BMC systems biology*, 4(1):1–14, 2010.
- [22] Roy S. Smith and John C. Doyle. Model invalidation: A connection between robust control and identification. In *1989 American Control Conference*, pages 1435–1440, 1989.
- [23] Tau-Mu Yi, Maryam Fazel, Xin Liu, Tosin Otitoju, Jorge Goncalves, Antonis Papachristodoulou, Stephen Prajna, and John Doyle. Application of robust model validation using SOSTOOLS to the study of g-protein signaling in yeast. In *Proceedings of the First Conference on Foundations of Systems Biology in Engineering*, pages 133–136. IFAC, 2005.

APPENDIX

A. Proof of Lemma 3.1

Proof: We begin showing that the extreme points of H are $\{\mathbf{v}_{ij} | i, j \in \mathbb{N}\}$. One can show this by considering the intersection of an arbitrary face of $\{\mathbf{p} | \|\mathbf{p}\|_1 \leq 2\eta\}$ and $\{\mathbf{p} | \sum_i p_i = 0\}$. Let F_α for $\alpha_i \in \{-2\eta, +2\eta\}$ be the face of $\{\mathbf{p} | \|\mathbf{p}\|_1 \leq 2\eta\}$ that is the convex hull of $\{\alpha_i \delta_i | i \in \mathbb{N}\}$. Since $F_\alpha = \{\sum_i \tau_i \alpha_i \delta_i | \sum_i \tau_i = 1, \tau_i \geq 0\}$, we have that $F_\alpha \cap \{\mathbf{p} | \sum_i p_i = 0\}$ is represented in the τ coordinates as $\tau \geq 0$ such that

$$\begin{bmatrix} 1 & 1 & 1 & \dots \\ \alpha_1 & \alpha_2 & \alpha_3 & \dots \end{bmatrix} \tau = \begin{bmatrix} 1 \\ 0 \end{bmatrix}.$$

One can see that the extreme points of this set are $\{\tau | \tau_i = 1/2, \tau_j = 1/2, \text{ for some } i, j \in \mathbb{N}\}$. Thus, the extreme points of $F_\alpha \cap \{\mathbf{p} | \sum_i p_i = 0\}$ are of the the form $\mathbf{p} = \eta \delta_i - \eta \delta_j$. Since for all $i, j \in \mathbb{N}$ there is an α such that $\mathbf{p} = \eta \delta_i - \eta \delta_j$ is an extreme point of $F_\alpha \cap \{\mathbf{p} | \sum_i p_i = 0\}$, the set of extreme points of H is $\{\eta \delta_i - \eta \delta_j | i, j \in \mathbb{N}\}$. Let

$$\begin{aligned} e &= \sup_{\alpha_{ij} \geq 0, \sum_{i,j=1}^{\infty} \alpha_{ij} = 1} \left\| K \sum_{i,j=1}^{\infty} \alpha_{ij} \mathbf{v}_{ij} + \mathbf{b} \right\|_2, \\ &= \sup_{\alpha_{ij} \geq 0, \sum_{i,j=1}^{\infty} \alpha_{ij} = 1} \left\| \sum_{i,j=1}^{\infty} \alpha_{ij} (K \mathbf{v}_{ij} + \mathbf{b}) \right\|_2, \\ &\leq \sup_{\alpha_{ij} \geq 0, \sum_{i,j=1}^{\infty} \alpha_{ij} = 1} \sum_{i,j=1}^{\infty} \alpha_{ij} \|K \mathbf{v}_{ij} + \mathbf{b}\|_2, \\ &\leq \sup_{0 < i, j} \|K \mathbf{v}_{ij} + \mathbf{b}\|_2. \end{aligned}$$

On the other hand, since $\{\mathbf{v}_{ij} | i, j \in \mathbb{Z}_{>0}\} \subset \{\mathbf{p} \in \ell_1 | \|\mathbf{p}\|_1 \leq 2\eta, \sum_i p_i = 0\}$, we have that $e \geq \sup_{0 < i, j} \|K \mathbf{v}_{ij} + \mathbf{b}\|_2$. Thus, $e = \sup_{0 < i, j} \|K \mathbf{v}_{ij} + \mathbf{b}\|_2$ and the proof is complete. ■

B. Proofs of Lemmas 3.2 and 3.3

To begin we state and prove a lemma that shows we can bound $\sup_{\mathbf{x} \in \mathbb{Z}_{\geq 0}^n : \mathbf{x} \notin N'} \|g^k(\mathbf{x})\|_2$ by $\max_{\mathbf{x} \in B(N')} \|g^k(\mathbf{x})\|_2$ if N' is sufficiently large.

Lemma B.1: Let $N' > k/\gamma$. We have that

$$\sup_{\mathbf{x} \in \mathbb{Z}_{\geq 0}^n: \mathbf{x} \not\leq N'} \|\mathbf{g}^k(\mathbf{x})\|_2 \leq \max_{\mathbf{x} \in B(N')} \|\mathbf{g}^k(\mathbf{x})\|_2.$$

Proof: Let $\mathbf{x}^* \not\leq N'$ and define $\mathbf{z} \in \mathbb{Z}_{\geq 0}^n$ as $z_i = \min\{x_i^*, N'\}$. Since $\mathbf{x} \not\leq N'$, there exists i such that $x_i^* > N'$, and hence $z_i = N'$. Thus, $\mathbf{z} \in B(N')$. Consider $g_i(\mathbf{x})$, an arbitrary element of $\mathbf{g}^k(\mathbf{x})$. We have that

$$\frac{\partial g_i(\mathbf{z} + \alpha(\mathbf{x}^* - \mathbf{z}))}{\partial \alpha} = \sum_{j=1}^n \frac{\partial g_i}{\partial x_j}(\mathbf{z} + \alpha(\mathbf{x}^* - \mathbf{z}))(x_j^* - z_j).$$

For each $j = 1, 2, \dots, n$ we have that either $x_j^* - z_j = 0$, or $z_j = N'$ and $x_j^* > N'$. Let us consider a j where the later is true.

$$\frac{\partial}{\partial x_j} g_i(\mathbf{x}) = e^{-\gamma(1, \mathbf{x})} \prod_{l \neq j} (x_l)^{i_l} \left(-\gamma (x_j)^{i_j} + i_j (x_j)^{i_j-1} \right),$$

which is negative as long as $x_j > i_j/\gamma$. For $0 \leq \alpha \leq 1$ we have $z_j + \alpha(x_j^* - z_j) \geq z_j \geq N'$. By assumption, $N' > k/\gamma$, so since $k \geq i_j$, we have that for all $0 \leq \alpha \leq 1$, $z_j + \alpha(x_j^* - z_j) > i_j/\gamma$ and hence

$$\frac{\partial g_i}{\partial x_j}(\mathbf{z} + \alpha(\mathbf{x}^* - \mathbf{z})) < 0.$$

Therefore, for all $0 \leq \alpha \leq 1$ and $i \leq k$,

$$\frac{\partial g_i(\mathbf{z} + \alpha(\mathbf{x}^* - \mathbf{z}))}{\partial \alpha} < 0,$$

and hence $\mathbf{g}^k(\mathbf{x}^*) < \mathbf{g}^k(\mathbf{z})$. We have show that for all $\mathbf{x}' \not\leq N'$, there exists a $\mathbf{z} \in B(N')$ such that $\|\mathbf{g}^k(\mathbf{x}')\|_2 \leq \|\mathbf{g}^k(\mathbf{z})\|_2$, and so

$$\max_{\mathbf{x}' \in \mathbb{Z}_{\geq 0}^n: \mathbf{x}' \not\leq N'} \|\mathbf{g}^k(\mathbf{x}')\|_2 \leq \max_{\mathbf{x} \in B(N')} \|\mathbf{g}^k(\mathbf{x})\|_2,$$

which completes the proof. \blacksquare

We now prove Lemmas 3.2 and 3.3.

Proof: [Lemma 3.2] To prove the first statement it suffices to show that $e_\lambda^2 \leq \tilde{e}_\lambda^2$ and $e_\lambda^3 \leq \tilde{e}_\lambda^3$. We have that

$$\begin{aligned} e_\lambda^2 &\leq \eta \cdot \sup_{(\mathbf{x}, \mathbf{x}') \in S_{N'}^2} (\|\mathbf{g}^k(\mathbf{x})\|_2 + \|\mathbf{g}^k(\mathbf{x}')\|_2), \\ &= 2\eta \cdot \max_{\mathbf{x} \in \mathbb{Z}_{\geq 0}^n: \mathbf{x} \not\leq N'} \|\mathbf{g}^k(\mathbf{x})\|_2. \end{aligned}$$

and thus by Lemma B.1, $e_\lambda^2 \leq \tilde{e}_\lambda^2$. We now show that $e_\lambda^3 \leq \tilde{e}_\lambda^3$.

$$\begin{aligned} e_\lambda^3 &\leq \eta \cdot \sup_{(\mathbf{x}, \mathbf{x}') \in S_{N'}^3} (\|\mathbf{g}^k(\mathbf{x})\|_2 + \|\mathbf{g}^k(\mathbf{x}')\|_2), \\ &\leq \eta \left(\max_{\mathbf{x} \in \mathbb{Z}_{\geq 0}^n: \mathbf{x} \leq N'} \|\mathbf{g}^k(\mathbf{x})\|_2 + \max_{\mathbf{x}' \in \mathbb{Z}_{\geq 0}^n: \mathbf{x}' \not\leq N'} \|\mathbf{g}^k(\mathbf{x}')\|_2 \right). \end{aligned}$$

It then follows from Lemma B.1 that $e_\lambda^3 \leq \tilde{e}_\lambda^3$, establishing that for all $N' > k/\gamma$, $e_\lambda \leq \tilde{e}_\lambda(N')$. We now show that $\lim_{N' \rightarrow \infty} \tilde{e}_\lambda(N') = e_\lambda$. First, observe that $\lim_{N' \rightarrow \infty} \tilde{e}_\lambda^1(N') = e_\lambda$. We will establish that asymptotically \tilde{e}_λ^2 and \tilde{e}_λ^3 are no larger than e_λ . For any i we have that $\lim_{\mathbf{x} \rightarrow \infty} g_i(\mathbf{x}) = 0$, and

hence $\lim_{N' \rightarrow \infty} \max_{\mathbf{x} \in B_{N'}} \|\mathbf{g}^k(\mathbf{x})\|_2 = 0$. Therefore, $\lim_{N' \rightarrow \infty} \tilde{e}_\lambda^2 = 0 \leq e_\lambda$. Additionally, observe that $\lim_{N' \rightarrow \infty} \tilde{e}_\lambda^3 = \eta \cdot \sup_{\mathbf{x} \in \mathbb{Z}_{\geq 0}^n} \|\mathbf{g}^k(\mathbf{x})\|_2 \leq e_\lambda$. Finally, $\lim_{N' \rightarrow \infty} \tilde{e}_\lambda(N') = e_\lambda$. \blacksquare

Proof: [Lemma 3.3] The proof proceeds similarly to the proof of Lemma 3.2. To prove the first statement we show that $e_T^2 \leq \tilde{e}_T^2(N')$ and $e_T^3 \leq \tilde{e}_T^3(N')$.

$$\begin{aligned} e_T^2 &\leq \sup_{(\mathbf{x}, \mathbf{x}') \in S_{N'}^2} \|T\boldsymbol{\lambda}^k(\hat{\pi})\|_2 + \eta(\|T\mathbf{g}^k(\mathbf{x})\|_2 + \|T\mathbf{g}^k(\mathbf{x}')\|_2), \\ &\leq \|T\boldsymbol{\lambda}^k(\hat{\pi})\|_2 + 2\eta \sup_{\mathbf{x} \in \mathbb{Z}_{\geq 0}^n: \mathbf{x} \not\leq N'} \|T\mathbf{g}^k(\mathbf{x})\|_2 \end{aligned}$$

For $N' > k/\gamma$ it then follows from Lemma B.1 that $e_T^2 \leq \tilde{e}_T^2$. Turning our attention to e_T^3 , let

$$a = \sup_{\substack{\mathbf{x}, \mathbf{x}' \in \mathbb{Z}_{\geq 0}^n: \\ \mathbf{x} \leq N', \mathbf{x}' \not\leq N'}} \|T\boldsymbol{\lambda}^k(\hat{\pi}) + \eta T\mathbf{g}^k(\mathbf{x}) - \eta T\mathbf{g}^k(\mathbf{x}')\|_2$$

and

$$b = \sup_{\substack{\mathbf{x}, \mathbf{x}' \in \mathbb{Z}_{\geq 0}^n: \\ \mathbf{x} \leq N', \mathbf{x}' \not\leq N'}} \|T\boldsymbol{\lambda}^k(\hat{\pi}) - \eta T\mathbf{g}^k(\mathbf{x}) + \eta T\mathbf{g}^k(\mathbf{x}')\|_2$$

It is true that $e_T^3 = \max\{a, b\}$. We have that

$$\begin{aligned} a &\leq \max_{\mathbf{x} \in \mathbb{Z}_{\geq 0}^n: \mathbf{x} \leq N'} \|T\boldsymbol{\lambda}^k(\hat{\pi}) + \eta T\mathbf{g}^k(\mathbf{x})\|_2 \\ &\quad + \eta \cdot \sup_{\mathbf{x}' \in \mathbb{Z}_{\geq 0}^n: \mathbf{x}' \not\leq N'} \|T\mathbf{g}^k(\mathbf{x}')\|_2, \end{aligned}$$

and hence by Lemma B.1 we have that

$$\begin{aligned} a &\leq \max_{\mathbf{x} \in \mathbb{Z}_{\geq 0}^n: \mathbf{x} \leq N'} \|T\boldsymbol{\lambda}^k(\hat{\pi}) + \eta T\mathbf{g}^k(\mathbf{x})\|_2 \\ &\quad + \eta \cdot \max_{\mathbf{x}' \in B_{N'}} \|T\mathbf{g}^k(\mathbf{x}')\|_2. \quad (12) \end{aligned}$$

By an analogous argument, one can show that

$$\begin{aligned} b &\leq \max_{\mathbf{x} \in \mathbb{Z}_{\geq 0}^n: \mathbf{x} \leq N'} \|T\boldsymbol{\lambda}^k(\hat{\pi}) - \eta T\mathbf{g}^k(\mathbf{x})\|_2 \\ &\quad + \eta \cdot \max_{\mathbf{x}' \in B_{N'}} \|T\mathbf{g}^k(\mathbf{x}')\|_2. \quad (13) \end{aligned}$$

Defining the right hand side of (12) as \tilde{a} and the right hand side of (13) as \tilde{b} , we have that $\tilde{e}_T^3 = \max\{\tilde{a}, \tilde{b}\}$, and since $a \leq \tilde{a}$ and $b \leq \tilde{b}$, we have that $e_T^3 \leq \tilde{e}_T^3$ for all $N' > k/\gamma$. To prove the second statement, we observe that $\lim_{N' \rightarrow \infty} \max_{\mathbf{x} \in B_{N'}} \|T\mathbf{g}^k(\mathbf{x})\|_2 = 0$, and hence $\lim_{N' \rightarrow \infty} \tilde{e}_T^2 \leq e_T$. Additionally, $\lim_{N' \rightarrow \infty} \tilde{e}_T^3 \leq e_T$. Therefore, $\lim_{N' \rightarrow \infty} \tilde{e}_T = e_T$. \blacksquare