

Keeping Your Story Straight: Truth-telling and Liespotting

Johannes Hörner*, Xiaosheng Mu†, Nicolas Vieille‡

December 18, 2017

Abstract

An agent privately observes a Markov chain online and reports it to a designer. To what patterns in the reported data should the designer pay attention? We show that, in general, keeping track of the empirical frequency of transition counts in the agent’s reports is insufficient, despite the true state being Markovian. Nonetheless, we derive conditions under which any deviation that can be distinguished from truth-telling by checking the frequency of strings of an arbitrary (finite) size can be detected by “checking pairs.” Further, we find that some undetectable deviations cannot be profitable, independent of the agent’s preferences. Hence, we provide weaker sufficient conditions that ensure that the agent finds honesty to be the best strategy. We explore the implications of these results for the literature on (i) linking incentives, (ii) dynamic implementation, and (iii) repeated games and agency models.

Keywords: Detectability, testing, Markov chains, implementation, repeated games.

JEL codes: C72, C73

*Yale University, 30 Hillhouse Ave., New Haven, CT 06520, USA, and Toulouse School of Economics, 21 Allée de Brienne, 31015 Toulouse, France, joh.horner@gmail.com

†Harvard University, 1805 Cambridge St., Cambridge, MA 02138, indefatigablexs@gmail.com.

‡HEC, 1, rue de la Libération, 78351 Jouy en Josas, France, vieille@hec.fr.

1 Introduction

A liar should have a good memory.

—Quintilian

If you tell the truth, you don't have to remember anything.

—Mark Twain

The least initial deviation from the truth is multiplied later thousandfold.

—Aristotle

Time and repetition provide opportunities to screen. To do so well, one must interpret the evidence that accumulates over time. Persistence in private information makes this especially valuable, as it constrains the scope for lying. However, it also raises a challenge: what patterns in the reports should one look for? An agent observes a Markov chain online and reports the states as they arrive to a designer. Preferences are summarized by the agent's reward function over state and report; technology is summarized by the Markov matrix. The designer seeks to elicit truth-telling in every round. She merely decides which report sequences are acceptable. In particular, she can require that the agent's reports be consistent with specific zero-one laws.

For instance, the designer might constrain the agent to choosing from those infinite sequences that are consistent with the long-run frequency of each state. Indeed, testing the state frequency is common in the literature (see below). It is both a simple and a surprisingly powerful test. Yet, because the underlying process is Markovian, it is natural to check for the empirical frequency of *pairs* of consecutive reports, requiring that this frequency matches the theoretical transition count of states. One might hope that such a test would outperform the former and to be without loss of generality. Our first example shows that it performs better, but it does not perform best. Prevarication adapts to liespotting: when pairs are checked, the agent gains from employing more sophisticated strategies, which satisfy the test of pairs but pay no heed to richer statistics.

We allow the designer to test for the frequency of k -tuples of reports (“testing k -tuples”), for any integer k , and address the following three questions.

First, what are the undetectable distributions (over states and reports) that the agent can engineer, for a given test? Second, when is truth-telling attainable, given the agent's preferences? These first two questions admit simple answers: the set of such distributions is a convex polytope,

with a distinguished vertex associated with truth-telling. Checking whether truth-telling is attainable with a given test amounts to checking whether the agent’s expected utility is maximized by this distinguished vertex—a simple matter of linear programming.

Third, for each of these two questions, is it without loss to consider states, pairs, or more generally, k -tuples, in the sense that considering longer strings will not help in reducing the set of undetectable distributions or in inducing truth-telling, independent of the agent’s preferences?

We show that testing (singleton) states is without loss of generality in some cases that are both special yet commonly assumed in economics: for instance, when there are two possible states only, or when the Markov chain is renewal. However, we prove that, provided that the Markov chain has three states or more, testing for singleton states only is with loss for almost all transition matrices. However, as suggested above, testing pairs is also not without loss, perhaps surprisingly. For some Markov chains, testing triples is better: doing so further reduces the set of undetectable distributions and expands the set of preferences for which speaking the truth is optimal.

There is nothing special about pairs. The same holds for k -tuples ($k \geq 2$). For some Markov chains, testing $k + 1$ -tuples shrinks the set of undetectable distributions that the agent can engineer and expands the set of preferences for which honesty is best. We obtain two further unexpected findings. First, while testing pairs (or k -tuples) is generally not without loss, there is an open positive-measure set of Markov matrices for which it is—unlike for singleton states, unless the state space is binary. Second, there is an open, positive-measure set of Markov matrices for which increasing k affects the set of undetectable distributions *yet* does not affect the range of preferences for which honesty is best. For such Markov chains, refining the test reduces the scope for lying, but the lies that get pruned out are irrelevant, independent of the agent’s preferences.

Our focus is on how the designer uses the information available to her, not on the instrument at her disposal. Yet, this instrument matters for the choice of relevant information, whether this is authority, as in our baseline model, or money, as in Section 4.1. There, we establish a duality between money and authority: restricting the agent to sequences that fit the theoretical frequency of k -tuples is equivalent to using transfers that only depend on the last k reports.¹ For instance, forcing the agent to match the theoretical frequency of each reported state amounts to assuming that the designer is only able to condition transfers on the latest report. Hence, money buys memory (and time: there is no need to wait literally forever to pass judgment).

¹Transfers with bounded memory satisfy desirable properties in related contexts. See, for instance, Bhaskar, Mailath and Morris (2013).

Our results have direct implications for implementation. In Section 4.2, we show the relationship between our results and the existing ones with respect to static implementation. We show how our results extend those of Rochet (1987) to a dynamic environment, couched in terms of cyclical monotonicity. By determining the set of incentive-feasible allocation functions, implementation is the first step in the design of the optimal mechanism given a particular objective function, such as efficiency or revenue maximization. Baron and Besanko (1984), Battaglini (2005), and other related works, are examples of such specific design problems, where the state (the valuation of the agent) follows a Markov chain. Examples of such hidden-knowledge agency problems arise frequently in regulation (*e.g.*, rate setting by public firms that have superior information regarding the current state of demand).

Our results have also immediate consequences for dynamic agency problems and Bayesian games. To the extent that there is no “universal” test—that is, no integer k that would be without loss, independent of the environment—there is also no hope of extending the celebrated “Bellman-type” characterizations of the equilibrium payoff set (whether for a given level of patience, as in Shapley, 1953, or Abreu, Pearce and Stacchetti, 1990, or for low discounting, as in Fudenberg and Levine, 1994), to environments with persistent information and interdependent values. As we explain in Section 4.3, keeping track of more complicated statistics of past evidence implies that, in general, the dynamic game cannot be summarized by a one-shot (or, more generally, k -shot) game.

Section 2 introduces an example that illustrates the main ideas. Section 3 generalizes the example and contains the main results. Section 4 develops the implications of the results regarding the role of money, implementation and dynamic games and agency.

Related Literature: Many papers have already examined the power of “linking incentive constraints” under imperfect observability. In the context of dynamic agency, early examples include Radner (1981), Townsend (1982) and Rubinstein and Yaari (1983). In the context of dynamic games, Abreu, Milgrom and Pearce (1991) show how tying together multiple draws help sustain cooperation (see also Samuelson and Stacchetti, 2016, for a more recent contribution). In the static context, Fang and Norman (2006), Jackson and Sonnenschein (2007) and Matsushima, Miyazaki and Yagi (2010) develop similar ideas. All these papers focus on the i.i.d. case and use a test that is a version of the frequency test for singleton states (variously referred to as “review strategies” or a “quota mechanism”). As Lemma 4 elucidates, the performance of a given test critically depends on the setup (dynamic vs. static), because of the agent’s information.

In a dynamic environment, the agent does not know future states in advance (even if, in the Markovian case, he has an informational advantage). As we show, if he were to know them, there is no need to go beyond singleton states.²

To the best of our knowledge, Escobar and Toikka (2013) are the first to consider a test that goes beyond testing singletons. Theirs is a variation on the test of pairs. As they argue in their Section 2, testing pairs improves on testing singletons in their Markovian environment. More broadly, our paper is closely related to the literature on repeated games with incomplete information that follows Aumann and Maschler (1995), in particular, Renault, Solan and Vieille (RSV, 2013) and Hörner, Takahashi and Vieille (2015). See also Athey and Bagwell (2008) and Barron (2017). There are two major differences between our work and these contributions. First, theirs are games without commitment, and hence, the solution concepts differ. Second, this literature focuses on the characterization of the set of equilibrium payoffs.³ In our environment, with a single agent, the range of possible payoffs is trivial: the agent’s payoff is minimized by a report-independent allocation function (which is clearly incentive-compatible), and the agent’s maximum payoff is certainly incentive-compatible. Nonetheless, some of the techniques and ideas are clearly related. In particular, the representation of undetectable report policies in terms of copulas when the state follows a pseudo-renewal chain is directly borrowed from RSV. In a dynamic mechanism design environment, Athey and Segal (2013) establish the existence of efficient, incentive-compatible and under certain assumptions budget-balanced mechanism. Here also, the focus on efficiency and private values implies that the mechanism can be represented in a recursive fashion.

Our paper yields a dynamic counterpart of Rochet (1987)’s characterization (see also Rockafellar, 1966; Afriat, 1967), who shows that cyclical monotonicity is a necessary and sufficient condition for implementation. Rahman (2010) provides a general (extensive-form) version of the result that implementation is equivalent to undetectable deviations being unprofitable. In our context, the entire difficulty is to determine which deviations affect the distribution over outcomes but are undetectable. Despite its name, the burgeoning literature on repeated implementation (see, for instance, Lee and Sabourian, 2011, 2015; Mezzetti and Renou, 2015; Renou and Tomala, 2016) is less related. In an environment without transferable utility, Lee and

²This does not mean that more complicated tests would not be useful in terms of rates of convergence, as shown by Cohn (2010) in the setup of Jackson and Sonnenschein (2007).

³Indeed, the main result of these papers is that under certain assumptions, the equilibrium payoff set can be described by a relatively simple recursive equation (involving, at most, pairs of states) or obtained by strategies that rely on testing at most pairs.

Sabourian and Mezzetti and Renou provide conditions on an allocation function in the spirit of Maskin monotonicity that guarantee (different versions of) implementability. These papers do not attempt a characterization in the Markovian case, which is our focus (although Renou and Tomala do allow for Markovian states).

We do not explore here what additional structure a specific objective function for the designer might impose. In the case of revenue maximization, Battaglini and Lamba (2015) show the difficulty of characterizing optimal mechanisms in Markovian environments with three or more states. Their analysis provides a nice contrast with Battaglini (2005), who focuses on two states only, and this difference resonates with the importance of memory when going to more than two states.

Our paper is also related to some literature in statistics. The problem of identifying hidden Markov chains has a long tradition in econometrics (see, among many others, Blackwell and Koopmans, 1957; Connault, 2016). Our problem is quite different, however, as there is no exogenous signal about the hidden Markov chain but instead a report, which is strategically and arbitrarily chosen as a function of the entire history, and hence, does not satisfy the usual assumptions imposed on the signal distribution.⁴ We cannot simply restrict the agent to using Markovian strategies: the more complicated the test applied to the agent is, the more sophisticated his best-reply.

Examples of non-Markovian processes that nevertheless satisfy the Chapman-Kolmogorov equation of a given (three-state or more) Markov chain were given by Feller (1959) and Rosenblatt (1960). However, simply because such examples exist does not imply that they matter for the joint distribution of states and reports, or for incentives.

2 An Example

Here, we develop an example to introduce the main ideas. We gloss over some of the formal or technical details, which are revisited in later sections.

In each round $n = 1, 2, \dots$, an agent (he) observes the realization of a Markov chain, with no

⁴See also Rothschild (1982), who show how partial observability turns a Markov system into a non-Markovian system.

foresight. The Markov chain takes values in $S = \{s_1, s_2, s_3\}$ according to the transition matrix

$$\begin{array}{c} s'_1 \quad s'_2 \quad s'_3 \\ \begin{array}{l} s_1 \\ s_2 \\ s_3 \end{array} \begin{pmatrix} 1/2 & 1/2 & 0 \\ 0 & 3/4 & 1/4 \\ 1/2 & 0 & 1/2 \end{pmatrix} \end{array}$$

That is, state s'_2 follows s_1 with probability (w.p.) $1/2$, etc. Plainly, the ergodic distribution λ assigns equal weight to s_1 and s_3 and more weight to s_2 : state 2 has invariant probability $1/2$, while states 1 and 3 have probability $1/4$ each.⁵

In round n , the agent is invited to make a report a_n regarding the prevailing state. For now, we assume that this report is an element from a copy of S , denoted A , to distinguish reports from states. Report a_i , $i = 1, 2, 3$, is interpreted as referring to state s_i .

The goal of the designer or principal (she) is to elicit truth-telling from the agent. Unfortunately, she has access to rather indigent information, as she observes the agent's reports only. Her instruments are equally primitive. Her only option is to impose constraints on the sequence of reports that the agent is allowed to produce. Neither the agent nor the designer are in a rush, and hence, that the fulfillment of these constraints can be evaluated "at the end of time."

For instance, she can require the agent to report s_2 one-half of the time and the other two states one-quarter of the time. An honest agent passes this test (a.s.). Indeed, a large literature has focused on (some version of) this test, showing how it implements any *ex ante* efficient social choice function (see Townsend (1982), Jackson and Sonnenschein (2007)). With this test, the designer is as effective as if she had transfers at her disposal, in a one-shot interaction.

However, our designer may not wish to maximize the agent's payoff only. For instance, she might take into account the welfare of other (uninformed and unmodeled) agents whose payoff also depends on the state. This might not suit the agent. Truth-telling might be beyond reach, whatever constraint is placed on the report sequence. The question, then, is the following: assuming that it is possible to keep the agent honest, what is a simple way to do so?

We bypass the underlying decision problem that maps reports into outcomes by positing some utility function for the agent over state and report, $r : S \times A \rightarrow \mathbf{R}$. His realized payoff over N

⁵While this is irrelevant, we may think of the initial state as being drawn according to this ergodic distribution.

rounds is then

$$\frac{1}{N} \sum_{n=1}^N r(s_n, a_n),$$

and his goal is to maximize the (lower) limit as $N \rightarrow \infty$ of his expected payoff over admissible sequences. Equivalently, the agent chooses a distribution $\mu \in \Delta(S \times A)$ over states and reports, interpreted as the limit frequency of these pairs arising under some reporting policy. Viewed in this way, the agent maximizes

$$\mathbf{E}_\mu[r(s, a)],$$

over those distributions μ that he can engineer through a reporting policy that abides by the constraints imposed by the designer. To be concrete, let us focus on a subset of such distributions. For $x \in [0, 1/4]$, consider the distribution μ_x over states and reports given by

$$\begin{array}{ccc} & a_1 & a_2 & a_3 \\ \begin{array}{l} s_1 \\ s_2 \\ s_3 \end{array} & \begin{pmatrix} \frac{1}{4} - x & x & 0 \\ 0 & \frac{1}{2} - x & x \\ x & 0 & \frac{1}{4} - x \end{pmatrix} \end{array}$$

When emulating such a distribution, the agent “exaggerates” the state a fraction $3x$ of the time, pretending in those rounds that the prevailing state is $s_{i+1} \pmod{3}$ when it is s_i . A high x would be attractive to the agent if his preferences were, say, $r(s_i, a_{i+d}) = d$, $i = 1, 2, 3$, $d = -1, 0, 1$. Note that telling the truth corresponds to the case in which $x = 0$, and we set $\mu^{tt} := \mu_0$.

What distributions μ_x can the agent bring about while passing the test described above? He can engineer *any* desired distribution in this class. Here is a strategy that does the trick: in s_i , he mixes between reporting a_i and a_{i+1} , independent of his private history, with suitable probabilities; *e.g.*, in s_1 , he reports a_1 w.p. $1 - 4x$, and a_2 w.p. $4x$. The key is that the marginal distribution of μ_x on A (the sum over each column) is equal to the invariant distribution λ . For each $\mu \in \mathcal{M}_0 = \{\mu \in \Delta(S \times A) : \text{marg}_S \mu = \text{marg}_A \mu = \lambda\}$, one can construct such a policy.

Yet it is easy to spot the agent’s lie when emulating $\mu_{\frac{1}{4}}$, that is, the distribution

$$\begin{array}{c} a_1 \quad a_2 \quad a_3 \\ s_1 \begin{pmatrix} 0 & \frac{1}{4} & 0 \end{pmatrix} \\ s_2 \begin{pmatrix} 0 & \frac{1}{4} & \frac{1}{4} \end{pmatrix} \\ s_3 \begin{pmatrix} \frac{1}{4} & 0 & 0 \end{pmatrix} \end{array}$$

Indeed, to achieve $\mu_{\frac{1}{4}}$, the agent has no choice but to report a_1 when the state is s_3 . However, state s_3 sometimes occurs after s_2 , when the agent is willing to report a_2 one-half of the time without yet knowing what the next state will be. Hence, the agent must sometimes report a_2 followed by a_1 , a sequence that cannot occur if he were to report states truthfully. Checking pairs of consecutive reports suffices to detect such a lie. Intuitively, because states are not i.i.d., it is useful to check the frequency of pairs (of states), rather than (singleton) states. As we will see, this intuition is essentially correct: checking states rather than pairs (or triples, etc.) suffices *only* for i.i.d. chains, or rather, for a slightly broader class of Markov chains (Theorem 1).

Because the underlying state follows a Markov chain, it is natural to keep track of the frequency of pairs. (This is what Escobar and Toikka (2013) do, for instance.) However, is it enough? Consider the distribution $\mu_{\frac{1}{6}}$ given by

$$\begin{array}{c} a_1 \quad a_2 \quad a_3 \\ s_1 \begin{pmatrix} 1/12 & 1/6 & 0 \end{pmatrix} \\ s_2 \begin{pmatrix} 0 & 1/3 & 1/6 \end{pmatrix} \\ s_3 \begin{pmatrix} 1/6 & 0 & 1/12 \end{pmatrix} \end{array}$$

Given a report a_1 or a_3 , the “lower” state (s_3 or s_2) occurs twice as often as the reported state.

Suppose that the previous report was a_3 and the current state is s_2 . What should the agent report? Reporting a_1 is not an option, as the entry (s_2, a_1) of $\mu_{\frac{1}{6}}$ is assigned probability 0. He also cannot report a_2 , as he reported a_3 previously: state s_3 cannot be followed by s_2 . Hence, he has no choice but to report a_3 again. This implies that he reports a_3 after a_3 at least one-half of the time, as (according to the joint distribution he emulates) when reporting a_3 , the state is s_2 two-thirds of the time (*i.e.*, $p(s_2|a_3) = 2/3$), and state s_2 repeats itself w.p. $3/4$; hence, state s_2 occurs after report a_3 at least $\frac{2}{3} \cdot \frac{3}{4} = \frac{1}{2}$ of the time. This is the frequency with which state

s_3 repeats itself; hence, after a_3 , the agent cannot report a_3 if any other state is realized, lest this be reflected by an excessive frequency of the pair of reports (a_3, a_3) . Hence, after a_3 , the agent reports a_3 if, and only if, the state is s_2 ; otherwise, he must report a_1 (as he cannot report a_2 after a_3). This is the unique reporting policy after a report a_3 that generates the correct frequency over pairs (a_3, a_3) and is consistent with the joint distribution that is being replicated.

Suppose, however, that our designer keeps track of *triples*. What if she observes two consecutive reports of a_3 ? Given the argument above, she can deduce that the current state must be s_2 , and thus, the next state must also be s_2 w.p. $3/4$, whereupon the agent must report a_3 again. Hence, she can predict that the next report will again be a_3 w.p. $3/4$, which contradicts the true frequency with which the state s_3 repeats itself: the policy of the agent is then distinguished from truth-telling, as must be any policy achieving the desired joint distribution.

The intuition is the following. To have the correct frequency of pairs, the agent must not only take into account his current state (here, s_2) when selecting the appropriate report but also consider yesterday's report (a_3). A liar must indeed take advantage of his memory, unlike an honest agent. As a result, today's state statistically depends on yesterday's report, even conditioning on the current report. However, today's state affects tomorrow's state and hence tomorrow's report. As a result, tomorrow's report and yesterday's report are not independent, conditional on today's report –a clear violation of the Markov property that can be detected by checking triples.⁶

For completeness, we now demonstrate that if the designer keeps track of pairs only, the agent gets away with the distribution $\mu_{\frac{1}{6}}$. As argued above, he would be able to generate the correct frequency over pairs (a_3, a_3) . What about pairs (a_1, a_1) and (a_2, a_2) ? Suppose that, conditional on a prior report a_1 , his reports are given by⁷

$$\begin{array}{c} a_1 \quad a_2 \\ s_1 \begin{pmatrix} 1/3 & 2/3 \end{pmatrix} \\ s_2 \begin{pmatrix} 0 & 1 \end{pmatrix} \\ s_3 \begin{pmatrix} 1 & 0 \end{pmatrix} \end{array}$$

⁶The argument thus far hinges on the fact that following a_3 , the agent reports a_3 if, and only if, the current state is s_2 . However, our conclusion does not rely on the if and only if. As long as s_2 is more persistent than s_3 and $p(s_2|s_3)$ is small enough, the probability of reporting a_3 following a_3 (conditional on the current state being s_2) is too high, and the contradiction follows. Hence, the example is not non-generic.

⁷In this matrix, rows are states, columns are reports, and entries are probabilities of the report given the state (and the previous report).

while, conditional on the previous report a_2 , his report follows the rule

$$\begin{array}{c} a_2 \quad a_3 \\ s_1 \begin{pmatrix} 1 & 0 \\ 7/8 & 1/8 \\ 0 & 1 \end{pmatrix} \\ s_2 \\ s_3 \end{array}$$

How likely is a_1 to be followed by a_1 ? Recall that, according to the joint distribution, when a_1 is reported, the state is either s_3 or s_1 , the former being twice as likely as the latter. In the former case, the next state is equally likely to be s_3 , which triggers a_1 given the candidate policy, or s_1 , which leads to a_1 w.p. $1/3$; in the latter case, only if the next state is also s_1 (which occurs w.p. $1/2$) does report a_1 occur (w.p. $1/3$). Overall, given a_1 , the next report is a_1 again w.p.

$$\frac{2}{3} \left(\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot \frac{1}{3} \right) + \frac{1}{3} \cdot \frac{1}{2} \cdot \frac{1}{3} = \frac{1}{2},$$

which is the correct frequency. What about the probability that a_2 is followed by another report a_2 ? A similar decomposition reveals that it is equal to

$$\underbrace{\frac{1}{3}}_{s^{-1}=s_1} \left(\underbrace{\frac{1}{2}}_{s=s_1|s^{-1}=s_1} \cdot \underbrace{1}_{a=a_2|s=s_1} + \underbrace{\frac{1}{2}}_{s=s_2|s^{-1}=s_1} \cdot \underbrace{\frac{7}{8}}_{a=a_2|s=s_2} \right) + \underbrace{\frac{2}{3}}_{s^{-1}=s_2} \cdot \underbrace{\frac{3}{4}}_{s=s_2|s^{-1}=s_2} \cdot \underbrace{\frac{7}{8}}_{a=a_2|s=s_2} = \frac{3}{4},$$

where s^{-1} is the past state, and underbraces indicate the event the probability of which is being given. Hence, this reporting policy yields the correct frequency over pairs of identical reports and, hence, (in this example) over all pairs of reports: it is indistinguishable from truth-telling if the designer checks pairs only. It is immediate to verify that it yields the desired joint distribution.

Two remarks are in order. First, the surprising result, if any, is not that some deviating reporting strategies become detectable once triples, rather than pairs, are scrutinized. This is almost trivial. Rather, it is that checking triples affects the set of joint distributions over states and reports that the agent can generate in an undetectable fashion. To appreciate the difference, note that, with an i.i.d. chain, any restriction on pairs prunes some reporting deviations (for instance, repeating in even rounds what was truthfully reported in the previous round), but no restriction limits the set of joint distributions that the agent can engineer.

Second, we have insisted that the agent tell the truth in all rounds. This is in contrast with many papers, which only require that this be the case in “most” rounds, with high enough probability –certainly a more reasonable demand if the horizon were finite, for instance. Yet our conclusion is robust, in the following sense. Given any $\mu \in \Delta(S \times A)$ close enough to $\mu_{\frac{1}{6}}$, by continuity, there is no reporting policy under which both the joint distribution of state and report is approximately equal to μ , and the frequency of *triples* of reports is approximately correct. Hence, checking triples improves on checking pairs, even if only approximate truth-telling is required.

A first conclusion of the example is that, from the perspective of a statistician who is interested in minimizing the scope of undetectable lies (joint distributions), keeping track of pairs improves on singletons (the agent can no longer engineer $\mu_{\frac{1}{4}}$), but it is no panacea, as checking triples is even better (the agent can no longer engineer $\mu_{\frac{1}{6}}$). In our example, it is not difficult to see that triples are also not the answer. Keeping track of quadruples further reduces the scope of lies, and so forth.

From the perspective of an economist concerned with enforcing truth-telling, the conclusion of this example is even bleaker. Note that the agent’s payoff is written as

$$\mathbf{E}_{\mu}[r(s, a)] = \mu_x \cdot r = \mu^{tt} \cdot r + xJ \cdot r,$$

where

$$J = \begin{pmatrix} -1 & 1 & 0 \\ 0 & -1 & 1 \\ 1 & 0 & -1 \end{pmatrix}$$

meaning that truth-telling is best within the class of distributions $\{\mu_x : x \leq 1/4\}$ if, and only if, $J \cdot r \leq 0$, a property that is independent of x . Hence, checking pairs rather than singletons does not enlarge the set of preferences for which truth-telling is optimal, within this class. This suggests that, due to the linearity of utility in the distribution, the set of undetectable deviations is not necessarily the appropriate benchmark for the economist. As we argue below, it is the cone spanned by this set, and the difference matters.⁸

⁸Our distinction between what is of interest to the statistician versus the economist makes for dramatic effect, but there are good reasons that an economist might also be interested in the set of undetectable distributions –in particular, if truth-telling is out of reach, it matters for the second-best solution. (See also Section 4.3.)

3 Main Results

This section generalizes the example, addressing the following questions:

1. What are the distributions over states and reports that the agent can engineer, when the designer checks singletons (pairs, triples, etc.)?
2. When is truth-telling implementable, given such a test, and *given* the agent's preferences?
3. In each case, what is the simplest test that is without loss?

For the sake of clarity, we retain the simplest setup and abstract for now from important but auxiliary issues (transfers, discounting, etc.), which are relegated to Sections 4 and 5.

3.1 Setup

We now begin with an arbitrary, time-homogeneous, irreducible and aperiodic Markov chain $(s_n)_{n \geq 1}$, taking values in the finite set S . Transition probabilities are denoted $p(s'|s)$, or $p_{ss'}$, and $P = (p(s'|s))_{s,s'}$ is the transition matrix. We let $\lambda \in \Delta(S)$ denote the invariant probability vector. In an abuse of notation, we also write $\lambda \in \Delta(S^{k+1})$, $k \in \mathbf{N}$, for the invariant probability vector of strings of length $k + 1$, namely,

$$\lambda(s^{-k}, \dots, s) = \lambda(s^{-k})p(s^{-(k-1)}|s^{-k}) \cdots p(s|s^{-1}).$$

Finally, we are given an initial distribution q over the initial state s_1 .

As before, an agent privately observes the Markov chain online and makes unverifiable reports to the designer. Reports are elements from A , a copy of S . Because preference is over the current state and report (see below), the agent does not benefit from conditioning reports on past realizations of the Markov chain. Hence, a (reporting) policy is a map $\sigma = (\sigma_n)_{n \geq 1}$, with $\sigma_n : A^{n-1} \times S \rightarrow \Delta(A)$, mapping the agent's past reports and the current state into a possibly random report. We denote by \mathbf{P}_q the law of the entire sequence (s_n) of states and by $\mathbf{P}_{q,\sigma}$ the law of the sequence $(s_n, a_n)_n$ under the policy σ .

We are given a utility function $r : S \times A \rightarrow \mathbf{R}$. The agent's realized payoff over N rounds is

$$\frac{1}{N} \sum_{n=1}^N r(s_n, a_n),$$

and the agent seeks to maximize the lower limit of its expectation under $\mathbf{P}_{q,\sigma}$ over policies σ .⁹

The designer wishes to incentivize truth-telling. That is, she wishes the agent to use σ^{tt} , the policy that always reports the correct state, independent of past reports. Under truth-telling, the agent’s expected payoff is simply

$$\mathbf{E}_{\mu^{tt}}[r(s, a)],$$

where $\mu^{tt} \in \Delta(S \times A)$ is the distribution under truth-telling (*i.e.*, $\mu^{tt}(s_i, a_i) = \lambda(s_i)$ for all $s_i \in S$).

To achieve her goal, the designer can require the agent’s reports to satisfy statistical constraints. Fix an integer k . For $n \geq k + 1$, let $f_n^k(a^{-k}, \dots, a)$ denote the empirical frequency over rounds $i \leq n$, in which the $k + 1$ most recent reports are (a^{-k}, \dots, a) . Under truth-telling, $(f_n^k)_n$ converges to λ as $n \rightarrow +\infty$, with probability 1. Hence, a natural way of using the limit statistics of reports is to check whether (f_n^k) converges to λ .

One can certainly think of other tests that are not subsumed by a frequency test on strings (for instance, one could sample at random a subsequence of the sequence of reports and scrutinize it in some way or run some automaton or program with a “forbidden” state on the sequence). Some of our results do not rely on our focus on strings. (For instance, Theorem 1 does not: no test improves on testing singletons if the chain is pseudo-renewal.) Others do not, although for any particular choice, one could presumably obtain results paralleling ours.

We denote by Σ_k the set of initial distributions q and reporting policies σ , such that $f_n^k \rightarrow \lambda$, $\mathbf{P}_{q,\sigma}$ -a.s., and write $V^N(q, \sigma)$ for the resulting expected average reward over the first N rounds.

Definition 1 *Truth-telling is k -limit optimal if*

$$\liminf_{N \rightarrow \infty} V^N(q, \sigma) \leq \mathbf{E}_{\mu^{tt}}[r(s, a)],$$

for every $(q, \sigma) \in \Sigma_k$. Truth-telling is limit optimal if it is k -limit optimal for some k .

Hence, truth-telling is k -limit optimal if it suffices to require the agent’s reports to satisfy the correct frequency over strings of length $k + 1$.

3.2 Undetectability

In the example in Section 2, we saw how the agent engineers any joint distribution $\mu \in \Delta(S \times A)$ that passes the test based on singleton states, provided that μ has the correct marginal

⁹As an alternative, slightly weaker criterion, the agent maximizes the expectation of the lower limit. Our propositions do not rely on this choice.

distribution on reports. Recall the following:

Definition 2 Let \mathcal{M}_0 be the set of distributions $\mu \in \Delta(S \times A)$ such that there exists a reporting policy $\sigma : S \rightarrow \Delta(A)$, with $\mu = \lambda \otimes \sigma$ being the induced joint distribution, satisfying $\text{marg}_A \mu = \lambda$.

Fix $k \geq 0$, and a stationary reporting policy $\sigma : A^k \times S \rightarrow \Delta(A)$ mapping the last k reports and the current state into a report (alongside some arbitrary specification for the first k rounds). For obvious reasons, we say that such a policy has *memory* k . Such a policy induces a Markov chain over the set $A^k \times S \times A$. This chain need not be irreducible and thus may admit a non-trivial set of invariant measures. The following generalizes the definition of \mathcal{M}_0 introduced above.

Definition 3 Fix $k \geq 0$. Let \mathcal{M}_k be the set of distributions $\mu \in \Delta(S \times A)$ such that there exists a reporting policy $\sigma : A^k \times S \rightarrow \Delta(A)$ and an invariant distribution $\nu \in \Delta(A^k \times S \times A)$ for σ such that the following two properties hold:

(i) the marginal of ν over $S \times A$ is equal to μ .

(ii) the marginal of ν over A^{k+1} is equal to λ , the invariant distribution of (s_{n-k}, \dots, s_n) .

Intuitively, this is the set of distributions that can be generated by a policy that is indistinguishable from truth-telling, when strings of length $k + 1$ are scrutinized. The set \mathcal{M}_k is the relevant object for the statistician, as it summarizes the joint distributions that cannot be detected given a test based on such strings. In the example in Section 2, $\mu_{\frac{1}{6}}$ is in \mathcal{M}_1 but not in \mathcal{M}_2 . Formally, given (q, σ) , we denote by $\mu_{q, \sigma}^N$ the expected empirical distribution of states and reports over the first N rounds.

Proposition 1 Fix any initial distribution $q \in \Delta(S)$. For every $\mu \in \mathcal{M}_k$, there exists σ such that $(q, \sigma) \in \Sigma_k$ and

$$\lim_{N \rightarrow \infty} \mu_{q, \sigma}^N = \mu. \quad (1)$$

Conversely, for every $(q, \sigma) \in \Sigma_k$, such that $\mu := \lim_{N \rightarrow \infty} \mu_{q, \sigma}^N$ exists, one has $\mu \in \mathcal{M}_k$.

We suspect that Proposition 1 holds unconditionally. For simplicity, we only prove it under an assumption (stated in Appendix ??) that covers full-support transition functions and all examples discussed in the paper.¹⁰

¹⁰Because Proposition 2 builds upon Proposition 1, this assumption is maintained there.

How does one compute \mathcal{M}_k ? Its definition invokes the existence of a policy with certain properties. To this end, fortunately, it suffices to consider policies with memory k . The next lemma provides an alternative characterization of \mathcal{M}_k that dispenses with this and clarifies its geometric structure.

Lemma 1 *For every k , \mathcal{M}_k is a convex polytope.¹¹ It is the set of distributions $\mu \in \Delta(S \times A)$ such that the following linear system in $\nu(a^{-k}, \dots, a^{-1}, s, a)$ has a solution:*

$$\sum_{a^{-k}, \dots, a^{-1}} \nu(a^{-k}, \dots, a^{-1}, s, a) = \mu(s, a). \quad (2)$$

$$\sum_s \nu(a^{-k}, \dots, a^{-1}, s, a) = \lambda(a^{-k})p(a^{-(k-1)}|a^{-k}) \dots p(a|a^{-1}). \quad (3)$$

$$\sum_a \nu(a^{-k}, \dots, a^{-1}, s, a) = \sum_{s^{-1}} \left(p(s|s^{-1}) \cdot \sum_{a^{-(k+1)}} \nu(a^{-(k+1)}, \dots, a^{-2}, s^{-1}, a^{-1}) \right). \quad (4)$$

$$\nu(a^{-k}, \dots, a^{-1}, s, a) \geq 0. \quad (5)$$

Despite Lemma 1, this polytope is complicated to describe directly. In the special case $|S| = 3$ (which we focus on below), there exists a direct description of the set \mathcal{M}_1 by means of “dual” constraints on μ .¹² Nonetheless, this set is “well-behaved,” as the next lemma states.

Lemma 2 *For all k , the polytope \mathcal{M}_k has dimension $(|S| - 1)^2$ and varies continuously with the transition matrix.*

Lemma 2 implies that the inequality $\mathcal{M}_1 \neq \mathcal{M}_2$ obtained in the example in Section 2 is robust to small changes in the transition matrix. In particular, the inequality still holds if all entries are strictly positive.

Plainly, the sequence $(\mathcal{M}_k)_{k \in \mathbf{N}}$ is nested. In Section 3.4, we ask whether it is eventually constant.

3.3 Unprofitable Deviations

Given Proposition 1, the following should come as no surprise.

¹¹See Gale (1960), for instance.

¹²That is, there exists a finite set of linear inequalities on $\mu(s, a)$ (without existential quantifiers) that characterizes \mathcal{M}_1 .

Proposition 2 *Truth-telling is k -limit optimal if, and only if,*

$$\mathbf{E}_\mu[r(s, a)] \leq \mathbf{E}_{\mu^{tt}}[r(s, a)] \text{ for all } \mu \in \mathcal{M}_k. \quad (6)$$

Thus far, we have postulated a perfectly patient agent, taking the limit of means as the payoff criterion. This has several benefits. First, rounds have equal importance in terms of payoffs, meaning that the only asymmetry between rounds is in terms of information: the agent knows earlier states, not later states. Second, it delivers particularly clean results, and the relevant sets \mathcal{M}_k have a nice structure, allowing for further investigation.

However, our results have counterparts for low discounting.¹³ The set \mathcal{M}_k as defined here still plays a role for implementation under discounting, as we show next. The initial distribution $q(s_1)$ is given and arbitrary. Given $\delta < 1$ and a reporting policy σ , we denote by $\mu_{q,\sigma}^\delta \in \Delta(S \times A)$ the expected, discounted distribution of the pairs (state, report). That is, for fixed $(s, a) \in S \times A$,

$$\mu_{q,\sigma}^\delta(s, a) := (1 - \delta) \mathbf{E}_{q,\sigma} \left[\sum_{n=1}^{\infty} \delta^{n-1} \mathbf{1}_{(s_n, a_n)=(s,a)} \right].$$

As a counterpart to Proposition 1, any limit point of $\mu_{q,\sigma}^\delta$ as $\delta \rightarrow 1$ belongs to \mathcal{M}_k (see the proof of Proposition 1). As a counterpart to Proposition 2, we have the following.

Proposition 3 *Assume that*

$$\mathbf{E}_\mu[r(s, \phi(a))] < \mathbf{E}_{\mu^{tt}}[r(s, \phi(a))] \text{ for all } \mu \in \mathcal{M}_k \setminus \{\mu^{tt}\}.$$

Then, for every $\varepsilon > 0$, there exists $\delta_0 < 1$ such that the following holds. For every $\delta > \delta_0$, there exists a transfer function $t_\delta : A^{k+1} \rightarrow \mathbf{R}$, such that

$$\|\mu_{q,\sigma}^\delta - \mu^{tt}\| < \varepsilon$$

for all optimal reporting policies σ in the δ -discounted Markov decision problem induced by t_δ .

In particular, the discounted frequency of rounds in which the agent reports truthfully exceeds $1 - \varepsilon$.

¹³See Frankel (2016) for an analysis of quota mechanisms, in the spirit of Jackson and Sonnenschein (2007), for the case of discounting.

However, the sets \mathcal{M}_k do not provide a tight characterization of truth-telling given a test for the following reason. It could be that $\mathcal{M}_k \neq \mathcal{M}_{k+1}$, while all rays spanned with vertices adjacent to μ^{tt} (a vertex that both sets have in common) are identical. That is, these polytopes might define the same cone pointed at μ^{tt} . See Figure 1 for an illustration. In that case, despite the proper set inclusion, one cannot find a utility function $r(\cdot, \cdot)$ for which truth-telling maximizes expected utility within \mathcal{M}_{k+1} but not within \mathcal{M}_k .

Such a situation appears improbable, but it is not impossible. In our leading example, $\mu_{\frac{1}{4}} \notin \mathcal{M}_1$, but $\mu_{\frac{1}{4}} \in \mathcal{C}_1$, as both $\mu_0 = \mu^{tt}$ and $\mu_{\frac{1}{6}}$ are in \mathcal{M}_1 , and we can write $\mu_{\frac{1}{4}} = \mu^{tt} + \frac{3}{2} \cdot (\mu_{\frac{1}{6}} - \mu^{tt})$. In fact, this is a more general feature of our example: as k increases, the set $\{\mu_x : \mu_x \in \mathcal{M}_k\}$ shrinks but never reduces to μ^{tt} , and thus, the ray it defines remains the same.¹⁴ This motivates the following definition.

Definition 4 *Given $k \geq 0$, let*

$$\mathcal{C}_k := \{\mu \in \Delta(S \times A) : \mu = \mu^{tt} + \alpha(\mu' - \mu^{tt}), \text{ for some } \alpha \geq 0, \mu' \in \mathcal{M}_k\}.$$

Plainly, $\mathcal{M}_k = \mathcal{M}_{k+1} \Rightarrow \mathcal{C}_k = \mathcal{C}_{k+1}$, and

$$\mathcal{C}_k \neq \mathcal{C}_{k+1} \Leftrightarrow \exists r \in \mathbf{R}^{|S| \times |A|} : \mathbf{E}_{\mu^{tt}}[r(s, a)] = \max_{\mu \in \mathcal{M}_{k+1}} \mathbf{E}_{\mu}[r(s, a)] < \max_{\mu \in \mathcal{M}_k} \mathbf{E}_{\mu}[r(s, a)],$$

by a standard separation argument. That is, if the cones differ, there are preferences such that truth-telling is $k + 1$ -limit optimal but not k -limit optimal. Hence, from the perspective of incentives, the cone \mathcal{C}_k is the proper object of study.

Unlike \mathcal{M}_k , the cone \mathcal{C}_k is not necessarily continuous in the transition matrix, at least when the limiting matrix has zero entries.

3.4 How Much Memory?

Depending on the agent's preferences, it can be impossible to induce the agent to tell the truth. Nevertheless, we may ask when the restriction to Σ_k is without loss, in the sense that truth-telling is limit optimal if and only if it is k -limit optimal, that is, $\mathcal{C}_k = \mathcal{C}_{k'}$ for all $k' > k$.

¹⁴We prove in Appendix F.3 that, for any k , there exists $x_k > 0$ such that $\mu_x \in \mathcal{M}_k$ whenever $x \in [0, x_k]$. To show that $x_k \rightarrow 0$, note that a report a_3 is followed by at least ℓ other reports a_3 w.p. $4x(3/4)^\ell$, which given x , exceeds the probability of such a sequence under truth-telling, $(1/2)^\ell$ whenever ℓ is large enough. Hence, given $x > 0$, there exists k such that $\mu_x \notin \mathcal{M}_\ell$ for all $\ell \geq k$.




Figure 1: An example in which $\mathcal{M}_k \neq \mathcal{M}_{k+1}$, yet $\mathcal{C}_k = \mathcal{C}_{k+1}$.

The same can be asked about the set \mathcal{M}_k , the set of undetectable distributions. These are properties of the Markov chain alone.

3.4.1 Testing Singleton States

We begin with the simplest test, namely, singleton states. We define a particular class of Markov chains.

Definition 5 *The chain (s_n) is pseudo-renewal if $s' \neq s \Rightarrow p(s'|s) = \alpha_{s'}$, for some $(\alpha_s)_{s \in S} \geq 0$.*

That is, a chain is pseudo-renewal if the probability of a change to a given state is independent of the initial state, provided that they are distinct. Constant chains are pseudo-renewal chains, and thus are i.i.d. processes. Pseudo-renewal chains are linear (although not necessarily convex) combinations thereof. When the chain belongs to one of these two special classes, it is intuitively clear that the problem reduces to the static one. The following theorem shows that pseudo-renewal chains characterize the reduction.

Theorem 1 *It holds that $\mathcal{C}_0 = \mathcal{C}_k$ for all k if, and only if, (s_n) is pseudo-renewal.*

Indeed, as shown in Renault, Solan and Vieille (2013), if (s_n) is pseudo-renewal, then for each $\mu \in \mathcal{M}_0$, there exists a reporting policy inducing μ , such that the distribution of the entire

sequence (a_n) is equal to that of the sequence (s_n) ; hence, no test, however sophisticated, can detect the deviation.¹⁵ Conversely, we show that if (s_n) is not pseudo-renewal, there exists $r(\cdot, \cdot)$ such that (6) holds for $k = 1$ but not for $k = 0$. For that preference, truth-telling is 1-limit, but not 0-limit, optimal.

How common are pseudo-renewal chains? In what follows, genericity statements are relative to the class of Markov chains of a given size $|S|$, identified with a subset of $\mathbf{R}^{|S| \times (|S|-1)}$, via the transition matrix.

Lemma 3 *Any Markov chain with $|S| = 2$ is pseudo-renewal. The set of pseudo-renewal chains is non-generic for $|S| \geq 3$.*

It is tempting to interpret Lemma 3 as saying that testing singletons is never robustly optimal. However, it is worth stressing that its conclusion does not only rely on the Markov chain, but on the information structure. The agent only finds about s_n in round n , not before. Knowing states in advance increases the number of incentive constraints (the agent must be incentivized to tell s_n honestly independent of his foreknowledge of future states). Because the agent can lie in more sophisticated ways, more complex statistical tests are called for, one might surmise. To the contrary, our next result, Lemma 4, establishes that all tests but the simplest one become ineffective, as the agent can foil all others. When the agent is *prophetic*, that is, when he knows the entire realized sequence (s_n) before being asked to make reports, the designer is left with no better option than checking state frequencies, independent of the Markov chain.¹⁶

Lemma 4¹⁷ *Suppose that the agent is prophetic. Then, truth-telling can be achieved if, and only if, it can be achieved when checking singleton states.*

That is, foresight suffices to nullify any potential benefit from exploiting the correlation of the states.¹⁸ Jackson and Sonnenschein (2007)'s test is not only natural but, in this context (with an infinite number of linked decisions), without loss of generality for implementation, even when states have a “Markovian” correlation structure.

¹⁵To be clear, their result is even stronger than $\mathcal{M}_0 = \mathcal{M}_k$ for all k because it covers arbitrary tests, regardless of whether they are based on strings.

¹⁶Matsushima, Miyazaki and Yagi (2010) have already shown that checking frequencies is without loss for static implementation with i.i.d. types. In fact, given Theorem 1, it does not matter that their framework is static.

¹⁷As the proof establishes, it suffices that the agent knows the next $|S| + 1$ states ahead of time.

¹⁸The role of correlation in our main results might bring to mind those of Crémer and McLean (1988), which suggest that correlated signals/states call for payments that depend on the entire vector of reports. The role of time, as elucidated by Lemma 4, is an important difference. Second, Crémer and McLean (1988) rely on the competition between the agents, while we have a unique agent.

3.4.2 Testing Pairs and More

Given Lemma 3, we turn our attention to $|S| \geq 3$. The case $|S| = 3$ is by far the simplest, and we focus on it.

We begin with the sets \mathcal{M}_k . Pairs do not always suffice, as our leading example showed. The next result generalizes.

Proposition 4 *Let $|S| = 3$. Fix a neighborhood of the transition matrix for (s_n) i.i.d. and uniform. Then, this neighborhood contains two open sets of transition matrices, one with $\mathcal{M}_1 = \mathcal{M}_k$ for all k and one with $\mathcal{M}_1 \neq \mathcal{M}_k$ for some k .*

Proposition 4 suggests that the relationship between the different sets \mathcal{M}_k relies on subtle details. There is no recursiveness in these sets either: for instance, there are cases in which $\mathcal{M}_1 = \mathcal{M}_2 \supsetneq \mathcal{M}_3$. Hence, just because testing some longer string does not help does not imply that no longer other string helps. See Appendix F.2 for such a transition matrix.

One may wonder about the focus on strings of *consecutive* reports. However, there are Markov chains (again, see Appendix F.2) for which testing the frequencies of (a^{-1}, a) , (a^{-2}, a) , and (a^{-2}, a^{-1}) (that is, pairs of consecutive and non-consecutive reports) leads to a strictly larger set of undetectable distributions than when testing triples (a^{-2}, a^{-1}, a) .¹⁹

We now turn to the sets \mathcal{C}_k , and argue that the distinction better undetectable and “unprofitable for all payoff functions” matters.

Proposition 5 *Let $|S| = 3$. For every $k \geq 1$, there exists an open set of transition matrices for which $\mathcal{M}_k \neq \mathcal{M}_\infty$, yet $\mathcal{C}_k = \mathcal{C}_\infty$.*

The example from Section 2 belongs to this class. If one insists on exact truth-telling, there is no point in looking beyond singleton states in this example. Yet, this example is special in the following sense. Focusing on unprofitability rather than undetectability does not mean that a finite memory suffices, in the following sense.

Proposition 6 *Let $|S| = 3$. For every $k \geq 1$, $\mathcal{C}_k \neq \mathcal{C}_{k+1}$ for an open set of transition matrices.*

¹⁹This is not to say that checking pairs of consecutive reports is necessarily the best test based on pairs. For instance, depending on the transition matrix, it might be better or worse to keep track of pairs with one round in between, as numerical examples show. That is, making the dependence on the Markov chain explicit, the set $\mathcal{M}_1(P)$ is sometimes but not always included in the set $\mathcal{M}_1(P^2)$. Hence, against a designer with bounded memory, more persistence does not necessarily harm the informed agent. This contrasts with the result of Peşki and Toikka (2017) in the related context of dynamic Bayesian games.

Proposition 6 implies that, in some cases, checking triples (resp., quadruples, etc.) allows truth-telling when pairs (triples, etc.) would not. However, this leaves open the possibility that for any given Markov chain (in a full measure set), some string length suffices.

It is worth stressing that Proposition 6 does not depend on exact truth-telling, as opposed to some weaker notion. To see this, consider a transition function for which \mathcal{C}_2 is a strict subset of \mathcal{C}_1 . Then, there exists a reward function r such that

$$\max_{\mu \in \mathcal{M}_2} \mathbf{E}_\mu[r(s, a)] = \mathbf{E}_{\mu^{tt}}[r(s, a)] < \max_{\mu \in \mathcal{M}_1} \mathbf{E}_\mu[r(s, a)]. \quad (7)$$

By the equality in (7), truth-telling can be obtained when triples are checked. However, consider *any* reporting policy $\tilde{\sigma}$ that reports truthfully most of the time, with high probability. The long-run payoff of the agent is then close to $\mathbf{E}_{\mu^{tt}}[r(s, a)]$. Let $\mu \in \mathcal{M}_1$ be a distribution such that $\mathbf{E}_{\mu^{tt}}[r(s, a)] < \mathbf{E}_\mu[r(s, a)]$ and σ a reporting policy associated with μ and thus indistinguishable from truth-telling when only pairs are checked. Then, some “perfectly” undetectable deviation improves upon any reporting policy that is approximately truthful. Thus truth-telling fails to be, even approximately, incentive-compatible when only pairs are checked.

This section has focused on the “negative” aspects of the problem: finite memory isn’t enough very generally, whether or not one is interested in detectability or profitability. Further, detectability and profitability are distinct problems.

However, there are interesting classes of Markov chains for which memory two, say, suffice for one or the other objective, as Propositions 5 and 6 prefigure. Unlike for the case of singleton states, when memory higher than two is considered, these classes are no longer non-generic. Identifying such classes is a difficult problem that is considered in a separate paper, see Hörner, Mu and Vieille (2017).

4 Implications and Applications

4.1 Money vs. Memory

In many economic problems, the designer has no direct control over the agent’s reports, but she can influence them via transfers. Here, let us assume that the designer can choose a sequence $(t_n)_n$ of transfers, measurable and bounded functions of the sequence of reports. For reasons that will become clear, we slightly change the environment and assume that the horizon is doubly

infinite. States in rounds $n \leq 0$ are “publicly observed.” However, s_n is privately observed by the agent in round $n \geq 1$, as before. The environment is otherwise unchanged, with the agent maximizing the expectation of

$$\liminf_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N (r(s_n, a_n) + t_n). \quad (8)$$

We stress that the designer can no longer impose a direct constraint on the sequence of reports.

A transfer function is a *transfer with memory k* , where $k \in \mathbf{N}$, if all t_n are functions of the last report a_n , in addition to the k most recent ones, a_{n-k}, \dots, a_{n-1} . It is *stationary* if this function is independent of calendar time. Because of time-invariance, we write t instead of t_n . The requirement that the transfer function has memory k is *a priori* restrictive: for instance, if the transfer function has memory 0, only the current report matters, and all statistical evidence is discarded. The designer’s problem becomes separable over time. Truth-telling, then, can only be incentivized if it is possible to do so in the one-shot interaction with transfers. One understanding of the result of Jackson and Sonnenschein (2007) is that keeping track of states is equivalent to using transfers in the one-shot game. When it comes to spotting lies, money substitutes for memory. This insight generalizes.

Proposition 7 *Truth-telling is k -limit optimal if, and only if, truth-telling is optimal for some stationary transfer with memory k .*

The reader might wonder about the importance of stationarity in the definition of transfers with a given memory size. In the undiscounted case, we do not know. However, if payoffs are discounted, we can show the following. Assuming direct mechanisms, the restriction to stationary transfers is without loss.

Lemma 5 *If truth-telling is implementable (under the discounted criterion) via some transfer with memory k , it can be achieved with a stationary transfer with memory k .*

The proof of Lemma 5 relies on the following mathematical fact. Let $(u_n)_{n \geq 1}$ be any bounded sequence of real numbers, with values in $[0, 1]$. For each $j \geq 1$, denote by $u^j = (u_{n+j-1})_{n \geq 1}$ the shifted sequence beginning with u_j . Regard each sequence u^j as a point in $[0, 1]^{\mathbf{N}}$, endowed with the product topology. Then, the closed convex hull of the set $\{u^j, j \geq 1\}$ contains a constant sequence. This latter fact is proved by a (infinite-dimensional) hyperplane separation theorem.

4.2 Implementation

Having the agent tell the truth is only useful to the extent that it has payoff-relevant consequences. Implementation takes as a primitive an arbitrary set Y and preferences over states and outcomes $y \in Y$, captured by some utility function $u : S \times Y \rightarrow \mathbf{R}$, and addresses the following question, in the standard static setup: for which maps $\phi : A \rightarrow Y$ can one find transfers $t : A \rightarrow \mathbf{R}$ such that truth-telling is optimal, given the payoff $u(s, \phi(a)) + t(a)$? Because of the revelation principle, the focus on direct truthful mechanisms is without loss.

The relationship between our problem and implementation becomes manifest once transfers are introduced, as in Section 4.1, and once we define $r(s, a) = u(s, \phi(a))$, for all $(s, a) \in S \times A$, given the map $\phi : A \rightarrow Y$.

The case in which transfers do not rely on past reports at all –mechanisms with memory 0, in the parlance of Section 4.1– has been extensively studied in prior work. For such mechanisms, the agent’s problem becomes separable over time. Hence, it reduces to the static problem, for which implementability has been characterized by Rochet (1987). Given $\phi : S \rightarrow Y$, u is *cyclically monotone* if, for every finite sequence $s_0, s_1, \dots, s_m = s_0$ of states, it holds that

$$\sum_{i=0}^{m-1} (u(s_i, \phi(s_{i+1})) - u(s_i, \phi(s_i))) \leq 0.$$

Rochet proves that ϕ is implementable (that is, there exist transfers $t : A \rightarrow \mathbf{R}$ such that truth-telling is optimal, with $A = S$) if and only if u is cyclically monotone. Note that this statement makes no reference to the distribution of the states: implementability does not depend on the prior.

Taking the dual of Rochet’s characterization immediately yields that ϕ is implementable via a stationary direct mechanism with memory 0 if and only if

$$\mathbf{E}_\mu[u(s, \phi(a))] \leq \mathbf{E}_{\mu^{tt}}[u(s, \phi(a))] \text{ for all } \mu \in \mathcal{M}_0,$$

which is precisely the special case $k = 0$ of Proposition 2.²⁰ As stated, this corollary suggests that implementability might depend on the prior λ , which defines \mathcal{M}_0 . This is however misleading. The defining property of \mathcal{M}_0 is that the marginals on each coordinate (s and a) coincide, not

²⁰In the static setup, this dual formulation appears well known, although it is difficult to attribute to a specific author.

what this common marginal is equal to.

More generally, we have the following as an immediate corollary of Proposition 2 and Proposition 7.

Corollary 1 *Let $\phi : S \rightarrow Y$. The map ϕ is implementable via transfers of memory k if, and only if,*

$$\mathbf{E}_\mu[u(s, \phi(a))] \leq \mathbf{E}_{\mu^{tt}}[u(s, \phi(a))] \text{ for all } \mu \in \mathcal{M}_k.$$

Corollary 1 extends to the case in which the choice function ϕ also exhibits memory, up to minor changes. Corollary 2 below is a straightforward variant, the proof of which we will omit. Fix a memory p for the map ϕ , which is now a function from S^{p+1} into Y .

Definition 6 *Given $k \geq p$, let \mathcal{M}_k^p be the set of distributions $\mu \in \Delta(A^p \times S \times A)$ such that there exists a reporting policy σ and an invariant distribution $\nu \in \Delta(A^k \times S \times A)$ for σ such that the following two properties hold:*

(i) *the marginal of ν over $A^p \times S \times A$ is equal to μ ;*

(ii) *the marginal of ν over A^{k+1} is equal to λ .*

Corollary 2 *Let $\phi : S^{p+1} \rightarrow Y$. The map ϕ is implementable with transfers of memory k if, and only if,*

$$\mathbf{E}_\mu[u(s, \phi(a^{-p}, \dots, a^{-1}, a))] \leq \mathbf{E}_{\mu^{tt}}[u(s, \phi(a^{-p}, \dots, a^{-1}, a))] \text{ for all } \mu \in \mathcal{M}_k^p.$$

There is also a primal version of Corollary 1, a generalization of cyclical monotonicity to the dynamic setup, which we only sketch, quoting results from the literature on MDPs. Let transfers $t : A^{k+1} \rightarrow \mathbf{R}$ be such that truth-telling is optimal in the induced undiscounted MDP over the state space $A^k \times S$, and denote by V the value of the MDP. Then, there exist so-called *relative values* $h : A^k \times S \rightarrow \mathbf{R}$, with the following properties : (i) for each $(\omega, s) \in A^k \times S$, one has

$$V + h(\omega, s) = \max_{a \in A} \left(u(s, \phi(a)) + t(\omega, a) + \sum_{s' \in S} p(s' | s) h((\omega^{+1}, a), s') \right) \quad (9)$$

(denoting by (ω^{+1}, a) the sequence obtained by dropping the first entry of ω and appending a); (ii) the maximum is achieved for $a = s$. Conversely, the existence of V and h such that (i) and

(ii) hold implies the optimality of truth-telling when transfers are set to t , and therefore, ϕ is implementable with transfers of memory k .

Consequently, if ϕ is implementable with transfers of memory k , then truth-telling is optimal in the *static* problem defined by the right-hand side of (9). Hence, the function

$$g_\omega(s, a) := u(s, \phi(a)) + \sum_{s' \in S} p(s' | s) h((\omega^{+1}, a), s')$$

is cyclically monotone for each $\omega \in A^k$.

Conversely, if for all ω , the function g_ω is cyclically monotone, then there exists $t(\omega, \cdot)$ such that for each s , truthful reporting $a = s$ achieves the maximum in (9). If, in addition, t can be chosen in such a way that the maximum is equal to $V + h(\omega, s)$ for some V , then truth-telling is also optimal in the MDP.

Because of this last condition, the existence of a function h such that all functions g_ω are cyclically monotone is not sufficient for implementability with transfers of memory k . In the dynamic case, unlike in the static one, the primal version is “self-referential,” as the right-hand side of (9) depends (via h) on the transfers, the existence of which is precisely in question.

We conclude this section with two lemmata concerning the restriction to direct mechanisms. Defining mechanisms generally entails a significant amount of notation, which is done in Appendix B.2. Roughly, a mechanism specifies a sequence of arbitrary message spaces, as well as allocation and transfer functions mapping sequences of messages into outcomes and payments. Neither the message space nor the functions need to be time-invariant. However it nonetheless could be that a (not necessarily stationary) optimal policy of the agent, given this mechanism, happens to induce a map from histories to outcomes such that only the current state matters, almost surely. If this is the case, we say that this map $\phi : S \rightarrow Y$ is implementable.

Lemma 6 *If ϕ is implementable, it is implementable via some direct mechanism.*

However, this is not to say that direct mechanisms are without loss once a restriction on the memory of the mechanism is imposed. A mechanism has *memory* $k \in \mathbf{N}$ if the allocation and transfer functions only depend on the current message and on the k previous ones. The following lemma shows that any mechanism with memory k can be emulated by a memory-1 mechanism, by appropriately expanding the message space.

Lemma 7 *If ϕ is implementable via a mechanism with memory k , ϕ is also implementable via a mechanism with memory 1.*

Neither lemma relies on the payoff criterion.

As an alternative to our focus on infinite-horizon, Markovian implementation problems, one might hope that long, but finite horizon version of the implementation problem would be more tractable. They are, but unfortunately, with a finite horizon, implementation is only possible if it is possible in the one-shot game. This extends to more sophisticated schemes that apply to the infinite-horizon case, in particular, review phases. That is, following Radner (1981), among others, it is customary in the literature to segment the infinite horizon into non-overlapping “phases” of $k + 1$ rounds and define transfers on each phase separately, such that the agent’s problem becomes separable across phases. Assume that reports are submitted in each round, but transfers are made every $k + 1$ rounds, as a function of the last $k + 1$ reports only. Say that ϕ is k -implementable in phases if there is $t : A^{k+1} \rightarrow \mathbf{R}$ such that truth-telling is optimal given transfers t .

Note that each such phase is independent of any other. Thus, in round $k + 1$, irrespective of the reports a_1, \dots, a_k submitted thus far, truthful reporting is optimal in the *static* problem when transfers are given by $t(a_1, \dots, a_k, \cdot)$, meaning that ϕ is implementable with transfers of memory 0. This proves Proposition 8 below.

Proposition 8 *Let $\phi : S \rightarrow Y$ be given. Then, ϕ is k -implementable in phases (with or without discounting) if and only if ϕ is implementable with transfers of memory 0.*

This shows that considering finite-horizon versions of the implementation problem (or using review phases more generally) is a poor way of understanding what happens with an infinite horizon: those versions collapse to the static problem.²¹ Equivalently, this says that requiring truth-telling in finite-horizon problems is extraordinarily demanding.

4.3 Recursive Representations

Most results in dynamic economics rely on a recursive representation of the payoff set. Promised utility is used as a state variable; then, a simple Bellman-type equation provides a characterization of the equilibrium payoff correspondence, as a function of this variable (and possibly other state variables). This representation holds quite generally in the context of agency

²¹Arguably, given the non-stationary nature of the finite-horizon problem, it might make sense to consider history-dependent allocation rules. Nevertheless, the final allocation rule must be implementable in the one-shot game.

models (see Spear and Srivastava (1987) and Thomas and Worrall (1990)), as well as in repeated games, as long as attention is restricted to public strategies (Shapley (1953) and Abreu, Pearce and Stacchetti (1990)).

An important implication of our results is that such a representation does not extend to the case of persistent private information. More precisely, with only one agent (as assumed thus far), the payoff set is still easy to compute (the agent can be expected to report truthfully when the designer's objective is to maximize his payoff; and to minimize his payoff, we might also assume that he does not report anything). However, given Proposition 6, this does not generalize to implementable choice functions. As a result, it does not generalize to the equilibrium payoff set either when there is more than one agent and values are interdependent.²² Interdependence arises naturally in repeated games or repeated agency problems in which actions are hidden (*e.g.*, in repeated insurance problems, as in Rubinstein and Yaari, 1983). This may be bad news for the tractability of dynamic models without transfers, but it is good news for the power of long-run relationships to support cooperation. This is because it implies that continuation play is not simply an imperfect substitute for transfers: in a Markovian environment, it can improve on transfers, to the extent that equilibrium payoffs can be achieved that would not be attainable in a static setup with contractual transfers.

The next example, building on the example in Section 2, formally makes this point. Consider a three-player infinitely repeated game. Player 1 is informed of the state, while players 2 and 3 are not. The state follows the Markov chain of the example. Player 1 has three actions $a^{(1)} \in \{1, 2, 3\}$, while player 2 has a binary action $a^{(2)} \in \{0, M\}$, taken immediately after player 1's (here, $M > 0$ is a large parameter to be determined). Player 3's action is also binary, $a^{(3)} \in \{0, M\}$, and taken immediately after player 2's. Let A denote the set of action profiles. Payoffs are

$$r^1(s_j, a) = a^{(2)} + \begin{cases} 0 & \text{if } a^{(1)} = j, \\ 1 & \text{if } a^{(1)} = j + 1, \\ -L & \text{if } a^{(1)} = j - 1, \end{cases}$$

²²In the case of independent private values, a recursive representation exists; see Hörner, Takahashi and Vieille (2015).

where indices are modulo 3, and

$$r^2(s_j, a) = -a^{(2)} - a^{(3)} + \begin{cases} 1 & \text{if } a^{(1)} = j, \\ -1 & \text{otherwise.} \end{cases}$$

Finally, $r^3(s_j, a) = 0$. Actions are observed, but payoffs are not. We note that player 2's action is *de facto* a transfer and thus is irrelevant for efficiency (*i.e.*, maximizing $r^1 + r^2 + r^3$), and that we can use player 3 (a dummy) to incentivize player 2. The challenge is to use the action of player 2 to induce player 1 to tell the truth: it is readily verified that the efficient choice is for player 1 to play $a^{(1)} = j$ in state s_j , "truth-telling."

It can be checked, using notations from Section 2, that $x_k := \max\{x : \mu_x \in \mathcal{M}_k\}$ is strictly positive and converges to zero as $k \rightarrow +\infty$ (see Appendix F.3). From the specification of payoffs, it also holds that for each k , there is some large parameter L such that the distribution $\mu \in \mathcal{M}_k$ that maximizes $\mu \cdot r^1$ is precisely μ_{x_k} . Consequently, there are games with joint equilibrium payoffs arbitrarily close to one.²³

Now fix $k \in \mathbf{N}$. Given a function $t : A^k \rightarrow \mathbf{R}^3$, define the k -repeated game $\Gamma^k(t)$ as follows. Player 1 privately observes online s_1, \dots, s_k , drawn according to the invariant distribution λ in the initial round and following the Markov chain afterwards. All players perfectly monitor actions. Player i 's expected payoff given strategy profile σ is

$$v^i(\sigma) = \mathbf{E}_{\lambda, \sigma} \left[\frac{1}{k} \sum_{n=1}^k r^i(s_n, a_n) + t^i(a_1, \dots, a_k) \right].$$

Let $E^{k,L,M}(t)$ denote the set of Nash equilibria of this game. Then, define the surplus as

$$S^{k,L,M} = \sup_{\substack{t: A^k \rightarrow \mathbf{R}^3 \\ \sigma \in E^{k,L,M}(t)}} \sum_i v^i(\sigma),$$

subject to the balanced-budget constraint $\sum_i t^i(a_1, \dots, a_k) \leq 0$ for all (a_1, \dots, a_k) .

²³Indeed, set k such that the joint payoff at μ_{x_k} is close to one; next set L such that $\mu \cdot r^1$ is maximized at μ_{x_k} . By the argument in the proof of Proposition 7, there are stationary transfers with memory k such that the optimal reporting policy of player 1 implements μ_{x_k} . Pick M sufficiently large to sustain these transfers as equilibrium strategies.

It then follows from our main results (and a continuity argument found in Appendix I) that

$$\sup_{L,M} S^{k,L,M} < 1.$$

That is, the highest surplus in the finitely repeated game augmented with (budget-balanced) transfers is bounded away from the highest surplus in the infinite-horizon game. Intuitively, player 1 has an incentive to deviate to a strategy inducing μ_{x_k} , which is unverifiable in the k -repeated game by Propositions 2 and 7. This deviation shifts joint payoffs away from efficiency.

Thus for each k , there are parameters L and M such that computing the set of equilibrium payoffs would require considering a reduced-form with more than k stages.²⁴ However, we suspect that L cannot be chosen uniformly in k ; hence, for a given payoff function, there might well be a representation of the equilibrium payoff set in terms of a k -shot game, for a suitably large k .

5 Concluding Comments

It would be of interest to assume that the agent does not directly observe the state. What happens if the agent's information itself is only a signal about the underlying state of nature? A major difficulty, then, is that the sequence of observations is not itself a Markov chain.

Related environments, such as continuous-time Markov chains, remain entirely unexplored: what is the meaning of a round in continuous time? Considering the uniformized discrete-time Markov chain does not seem to be as helpful here as in related problems because the exact time elapsed since the last reported switch is a piece of information that the designer might want to use. What if the agent observes a diffusion process, such as a Brownian motion?

From an economic perspective, it is desirable to go beyond the rather extreme objective to achieve truth-telling, and achieve it for all preferences. If truth-telling is not achievable, we may still ask how close one can get to it, and how memory affects this distance. At an abstract level, using the revelation principle, the agent would truthfully report a belief (rather than a state), coinciding with the Bayesian posterior that the designer forms upon hearing this

²⁴Clearly, our claim is not that strategies with bounded memory are with loss in repeated games (a well-known fact) but that the equilibrium payoff set in Bayesian games cannot be solved for by analyzing a one-shot game, regardless of whether this one-shot game is parametrized by state variables, as is done in APS for instance (in their language, actions supportable by $W = \mathbf{R}$ are those that can be supported in a one-shot game with transfers). We also do not claim that *no* one-shot game exists with the same equilibrium payoff set. We have fixed the payoff function and the Markov chain to be the same as in the infinite-horizon game. Otherwise, the question becomes meaningless, as we can always design the payoff function to deliver whichever set is to be replicated.

report. Given a test, one then computes the distance between the resulting empirical distribution over posterior beliefs and the true distribution of the state (this calls for preferences r over $S \times \Delta(S)$). Conversely, insisting on truth-telling, one might be interested in identifying the class of preferences for which a given test suffice, or in finding conditions to verify whether a given allocation function is implementable simplify, in the spirit of the Spence-Mirrlees condition, under which cyclical monotonicity reduces to standard monotonicity (see Rochet (1987)). In a dynamic setup, truth-telling depends both on the preferences and the Markov chain. Our dual approach is promising in this regard: Since the bi-dual is just the primal, the dual inequality constraints for \mathcal{M}_k characterize the extremal reward functions that are k -implementable, allowing a characterization of those preferences for which a given test suffices. Such an analysis is an important step toward bridging the gap between results on implementation, such as ours, and the literature on dynamic mechanism design.

References

- [1] Abreu, D., D. Pearce and P. Milgrom (1991). “Information and Timing in Repeated Partnerships,” *Econometrica*, **59**, 1713–1733.
- [2] Abreu, D., D. Pearce and E. Stacchetti (1990). “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica*, **58**, 1041–1063.
- [3] Afriat, S.N. (1966). “The Construction of Utility Functions from Expenditure Data,” *International Economic Review*, **8**, 67–77.
- [4] Athey, S. and K. Bagwell (2008). “Collusion with Persistent Shocks,” *Econometrica*, **76**, 493–540.
- [5] Athey, S. and I. Segal (2013). “An Efficient Dynamic Mechanism,” *Econometrica*, **81**, 2463–2485.
- [6] Aumann, R.J. and M. Maschler (1995). *Repeated Games with Incomplete Information*. Cambridge, MA: MIT Press.
- [7] Baron, D. and D. Besanko (1984). “Regulation and Information in a Continuing Relationship,” *Information Economics and Policy*, **1**, 447–470.

- [8] Barron, D. (2017). “Attaining Efficiency with Imperfect Public Monitoring and Markov Adverse Selection,” *Theoretical Economics*, forthcoming.
- [9] Battaglini, M. (2005). “Long-Term Contracting With Markovian Consumers,” *American Economic Review*, **95**, 637–658.
- [10] Battaglini, M. and R. Lamba (2015). “Optimal Dynamic Contracting: the First-Order Approach and Beyond,” working paper, Cornell University.
- [11] Bhaskar, V., G.J. Mailath and S. Morris (2013). “A Foundation for Markov Equilibria in Sequential Games with Finite Social Memory,” *Review of Economic Studies*, **80**, 925–948.
- [12] Billingsley, P. (1961). “Statistical Methods in Markov Chains,” *Annals of Mathematical Statistics*, **32**, 12–40.
- [13] Blackwell, D. and L. Koopmans (1957). “On the Identifiability Problem for Functions of Finite Markov Chains,” *Annals of Mathematical Statistics*, **28**, 1011–1015.
- [14] Cohn, Z. (2010). “A note on linked bargaining,” *Journal of Mathematical Economics*, **46**, 238–247.
- [15] Connault, B. (2016). “Hidden Rust Models,” working paper, University of Pennsylvania.
- [16] Crémer, J. and R.P. McLean (1988). “Full Extraction of the Surplus in Bayesian and Dominant policy Auctions,” *Econometrica*, **56**, 1247–1257.
- [17] Escobar, P. and J. Toikka (2013). “Efficiency in Games with Markovian Private Information,” *Econometrica*, **81**, 1887–1934.
- [18] Fang, H. and P. Norman (2006). “To Bundle or Not To Bundle,” *RAND Journal of Economics*, **37**, 946–963.
- [19] Feller, W. (1960). “Non-Markovian processes with the semi-group property,” *Annals of Mathematical Statistics*, **30**, 1252–1253.
- [20] Frankel, A. (2016). “Discounted quotas,” *Journal of Economic Theory*, **166**, 396–444.
- [21] Fudenberg, D. and D. Levine (1994). “Efficiency and Observability in Games with Long-Run and Short-Run Players,” *Journal of Economic Theory*, **62**, 103–135.

- [22] Fudenberg, D., D. Levine, and E. Maskin (1994). “The Folk Theorem with Imperfect Public Information,” *Econometrica*, **62**, 997–1040.
- [23] Gale, D. (1960). *The Theory of Linear Economic Models*, University Of Chicago Press.
- [24] Hörner, J., X. Mu and N. Vieille (2017). “TBD,” working paper, Harvard, HEC Paris and Yale.
- [25] Hörner, J., S. Takahashi and N. Vieille (2015). “Truthful Equilibria in Dynamic Bayesian Games,” *Econometrica*, **83**, 1795–1848.
- [26] Jackson, M.O. and H.F. Sonnenschein (2007). “Overcoming Incentive Constraints by Linking Decisions,” *Econometrica*, **75**, 241–258.
- [27] Judd, K.L., S. Yeltekin and J. Conklin (2003). “Computing supergame equilibria,” *Econometrica*, **71**, 1239–1254.
- [28] Lee, J. and H. Sabourian (2011). “Efficient Repeated Implementation,” *Econometrica*, **79**, 1967–1994.
- [29] Lee, J. and H. Sabourian (2015). “Complexity and repeated implementation,” *Journal of Economic Theory*, **158**, 259–292.
- [30] Matsushima, H., K. Miyazaki and N. Yagi (2010). “Role of linking mechanisms in multitask agency with hidden information,” *Journal of Economic Theory*, **145**, 2241–2259.
- [31] Mezzetti, C. and L. Renou (2015). “Repeated Nash Implementation,” *Theoretical Economics*, forthcoming.
- [32] Peşki, M. and J. Toikka (2017). “Value of Persistent Information,” working paper, MIT.
- [33] Radner, R. (1981). “Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship,” *Econometrica*, **49**, 1127–1148.
- [34] Rahman, D. (2010). “Dynamic Implementation,” working paper, University of Minnesota.
- [35] Renault, J., E. Solan and N. Vieille (2013). “Dynamic Sender-Receiver Games,” *Journal of Economic Theory*, **148**, 502–534.

- [36] Renou, L. and T. Tomala (2015). “Approximate Implementation In Markovian Environments,” *Journal of Economic Theory*, **159**, 401–442.
- [37] Rochet, J.-C. (1987). “A necessary and sufficient condition for rationalizability in a quasi-linear context,” *Journal of Mathematical Economics*, **16**, 191–200.
- [38] Rockafellar, T. (1966). “Characterization of the subdifferentials of convex functions,” *Pacific Journal of Mathematics*, **17**, 497–510.
- [39] Rosenblatt, M. (1960). “An aggregation problem for Markov chains,” in *Information and Decision Processes*, R.E. Machol, ed., McGraw-Hill, New York, 87–92.
- [40] Rothschild, M. (1982). “A Note on Partially Observed Markov Systems,” *Economics Letters*, **10**, 207–209.
- [41] Rubinstein, A. and M. Yaari (1983). “Repeated Insurance Contracts and Moral Hazard,” *Journal of Economic Theory*, **30**, 74–97.
- [42] Samuelson, L. and E. Stacchetti (2016). “Even Up: Maintaining Relationships,” working paper, Yale University.
- [43] Shapley, L.S. (1953). “Stochastic Games,” *Proceedings of the National Academy of Sciences of the U.S.A.*, **39**, 1095–1100.
- [44] Spear, S. and S. Srivastava (1987). “On Repeated Moral Hazard with Discounting,” *Review of Economics Studies*, **54**, 599–617.
- [45] Thomas, J. and T. Worrall (1990). “Income Fluctuation and Asymmetric Information: An Example of a Repeated Principal-Agent Problem,” *Journal of Economic Theory*, **51**, 367–390.
- [46] Townsend, R. (1982). “Optimal Multi-Period Contracts and the Gain from Enduring Relationships under Private Information,” *Journal of Political Economy*, **90**, 1166–1186.

The structure of the Appendix replicates that of the paper. The only exception is that the proofs of the statements in Section 3.4.2, being much longer, are given at the end, in the Supplementary Online Appendix.

A Undetectability and Unprofitable deviations

A.1 Undetectability

A.1.1 Proof of Proposition 1

As indicated in the main text, we prove Proposition 1 under an additional assumption, which we introduce next. A stationary reporting policy σ maps the set $A^k \times S$ of previous reports and of current states into $\Delta(A)$ and therefore induces a Markov chain over the set $A^k \times S \times A$.

Assumption A: If this Markov chain admits an invariant distribution ν that satisfies condition (ii) in Definition 3, then the chain has a unique recurrent set.

Assumption **A** is satisfied when transitions have full support.²⁵ It is also satisfied when $|S| = 3$, and the support of each $p(\cdot | s)$ ($s \in S$) has cardinality at least 2, as in all examples discussed in the paper.²⁶ Since such a stationary reporting policy σ relies on the k most recent reports, the reporting behavior in the first k rounds is ambiguous. In some cases, we will implicitly assume that σ draws fictitious reports for rounds $n \leq 0$ in some arbitrary way. In most cases however, the drawing process will matter. We will then allow for initial distributions that select both fictitious reports and an initial state. To avoid introducing additional notations, such distributions will also select the actual first report, and will be denoted by $\nu_0 \in \Delta(A^k \times S \times A)$ to distinguish them from initial distributions $q \in \Delta(S)$. For such extended initial distributions, the probability distribution $\mathbf{P}_{\nu_0, \sigma}$ is well-defined. With a slight abuse of notation, we will then write $(\nu_0, \sigma) \in \Sigma_k$ whenever $f_n^k \rightarrow \lambda$, $\mathbf{P}_{\nu_0, \sigma}$ -a.s.

We start with the direct implication in Proposition 1. Fix an initial distribution q , and $\mu \in \mathcal{M}_k$. Let a pair (σ, ν) be associated with μ , as in Definition 3. Since ν is invariant for σ , $\mu_{\nu, \sigma}^N$ is equal to μ for each N . By ergodicity, the empirical frequency over the first n rounds of any sequence $(a^{-k}, \dots, a^{-1}, s, a)$ converges to $\nu(a^{-k}, \dots, a^{-1}, s, a)$, $\mathbf{P}_{\nu, \sigma}$ - and $\mathbf{P}_{q, \sigma}$ -a.s. as $n \rightarrow +\infty$.²⁷ Taking marginals over A^{k+1} , this implies $f_n^k \rightarrow \lambda$, $\mathbf{P}_{q, \sigma}$ -a.s. Hence $(q, \sigma) \in \Sigma_k$. By ergodicity again, the difference $\mu_{q, \sigma}^N - \mu_{\nu, \sigma}^N$ vanishes as $N \rightarrow +\infty$. This concludes the first part of the proof.

We turn to the proof of the converse implication. We will prove a slightly stronger statement, which we state as a lemma because it is used elsewhere.

²⁵Then, the unique recurrent set is equal to the set of all $(\vec{a}, s, a) \in A^k \times S \times A$ such that $\sigma(a | \vec{a}, s) > 0$.

²⁶Indeed, assume there are two distinct recurrent sets R and $R' \subset A^k \times S \times A$. Assume in addition that, for some $\vec{a} \in A^k$, $a \in A$, there are $s, s' \in S$, with $(\vec{a}, s, a) \in R$ and $(\vec{a}, s', a) \in R'$. Given our assumption on p , there is $\bar{s} \in S$ with $p(\bar{s} | s) > 0$ and $p(\bar{s} | s') > 0$, hence the sequence (\vec{a}, a, \bar{s}) —with some abuse of notation—lies in both R and R' , which cannot be. Hence, the projections on $A^k \times A$ of the different recurrent sets are pairwise disjoint. This can be seen to stand in contradiction with (ii) in Definition 3.

²⁷For this statement, the drawing process of fictitious reports is irrelevant.

Lemma 8 *Let $(q, \sigma) \in \Sigma_k$. Then any limit point of $\mu_{q, \sigma}^\delta$ (see Section ??), as $\delta \rightarrow 1$, belongs to \mathcal{M}_k .*

Note that if $\mu := \lim_{N \rightarrow \infty} \mu_{q, \sigma}^N$ exists, then $\mu = \lim_{\delta \rightarrow 1} \mu_{q, \sigma}^\delta$, hence the converse implication in Proposition 1 follows. We introduce the set \mathcal{M} of distributions $\mu \in \Delta(S \times A)$ that are limit points of $\mu_{q, \sigma}^\delta$ for some $(q, \sigma) \in \Sigma_k$. We will prove that

$$\sup_{\mu \in \mathcal{M}} \alpha \cdot \mu \leq \sup_{\mu \in \mathcal{M}_k} \alpha \cdot \mu,$$

for each $\alpha \in \mathbf{R}^{S \times A}$. Since \mathcal{M}_k is a (compact) polytope, this will imply that $\mathcal{M} \subset \mathcal{M}_k$, as desired.

Thus, fix $\alpha \in \mathbf{R}^{S \times A}$, and $\bar{\mu} \in \mathcal{M}$. Let $(\bar{q}, \bar{\sigma}) \in \Sigma_k$ and $(\delta_n)_n$ be a sequence converging to one, such that $\bar{\mu} = \lim_{n \rightarrow \infty} \mu_{\bar{q}, \bar{\sigma}}^{\delta_n}$. For given $\delta < 1$ and $\eta > 0$, consider the following constrained optimization problem: find a (general) reporting policy σ and an extended initial distribution ν_0 that maximizes $\alpha \cdot \mu_{\nu_0, \sigma}^\delta$, among the pairs (ν_0, σ) such that $\|\lambda_{\nu_0, \sigma}^\delta - \lambda\| \leq \eta$.²⁸ By Chapter 3 in Altman (1999), there is an optimal pair $(\nu_{\delta, \eta}, \sigma_{\delta, \eta})$ with $\sigma_{\delta, \eta}$ a stationary reporting policy.

Since $(\bar{q}, \bar{\sigma}) \in \Sigma_k$, one has $f_n^k \rightarrow \lambda$, $\mathbf{P}_{\bar{q}, \bar{\sigma}}$ -a.s., and therefore, $f_\delta^k \rightarrow \lambda$ as well, as $\delta \rightarrow 1$. Taking expectations, this implies that $(\bar{q}, \bar{\sigma})$ is a feasible point in the constrained optimization problem provided δ is close enough to one.²⁹ That is, there exists $\delta(\eta) < 1$ such that

$$\delta > \delta(\eta) \Rightarrow \alpha \cdot \mu_{\bar{q}, \bar{\sigma}}^\delta \leq \alpha \cdot \mu_{\nu_{\delta, \eta}, \sigma_{\delta, \eta}}^\delta. \quad (10)$$

Consider now any subsequence of $(\delta_m)_m$, still denoted $(\delta_m)_m$, and any sequence $(\eta_m)_m$ converging to zero such that $\delta_m > \delta(\eta_m)$ for each m . Let $\sigma_m = \sigma_{\delta_m, \eta_m}$, and let $\nu_m = \nu_{\delta_m, \sigma_m} \in \Delta(A^k \times S \times A)$ denote discounted empirical frequencies. Passing to a further subsequence if necessary, we assume that σ_m and ν_m have limits, denoted σ and ν . It is a standard technical fact that ν is invariant for σ (see Lemma 9 below). Since the marginal $\lambda_{q_m, \sigma_m}^{\delta_m}$ of ν_m converges to λ , the marginal of ν on A^{k+1} is λ , hence the limit $\mu \in \mathcal{M}_k$.

Applying (10) with (δ_n, η_n) and letting $n \rightarrow +\infty$, one obtains $\alpha \cdot \bar{\mu} \leq \alpha \cdot \mu$, as desired.

Lemma 9 *Let $(\delta_m)_m \rightarrow 1$, let $(\sigma_m)_m$ be a sequence of stationary reporting policies with memory k converging (in the usual sense) to σ , and let (q_m) be a sequence of initial distributions. Let $\mathcal{H} \subset \Delta(A^k \times S \times A)$ denote the set of invariant measures for σ . Then all limit points of $(\nu_{q_m, \sigma_m}^{\delta_m})_m$ belong to \mathcal{H} .*

Proof. For $\psi, \psi' \in A^k \times S \times A$, we denote by $\nu_m(\psi, \psi')$ be the δ_m -discounted frequency of ψ' , when starting a Markov chain from ψ and using the stationary policy σ_m . Write $Q_m = (\nu_m(\psi, \psi'))_{\psi, \psi'}$ as the stochastic matrix of size $|A^k \times S \times A|$. Pick an arbitrary convergent subsequence of $(Q_m)_m$, still denoted $(Q_m)_m$, with limit Q .

Let P_{σ_m} denote the transition matrix of the Markov chain induced by σ_m , then we have the identity

$$Q_m = (1 - \delta_m) \sum_{i=0}^{+\infty} \delta_m^i P_{\sigma_m}^i = (1 - \delta_m)I + \delta_m P_{\sigma_m} Q_m,$$

²⁸ $\lambda^\delta(\dots)$ is defined to be the empirical frequency of report strings with length $k+1$.

²⁹Because $\bar{\sigma}$ ignores early fictitious reports, the pair $(\bar{q}, \bar{\sigma})$ can be identified with any pair $(\bar{\nu}, \bar{\sigma})$, as long as the marginal on S of $\bar{\nu}$ is \bar{q} .

so that

$$Q_m = (1 - \delta_m)(I - \delta_m P_{\sigma_m})^{-1}. \quad (11)$$

Rewriting (11) as $(1 - \delta_m)I = Q_m - \delta_m Q_m P_{\sigma_m}$ and letting $m \rightarrow \infty$, one gets $Q = Q P_\sigma$. In particular, each row of Q is an invariant (probability) measure for σ . Hence, all limit points of ν_m are invariant for σ , completing the proof. ■

A.1.2 Proof of Lemma 1

Generic elements of $A^k \times S \times A$ are denoted $(a^{-k}, \dots, a^{-1}, s, a)$, with the interpretation that s and a are the current state and report, while a^{-i} is the report that was submitted i rounds earlier. Given $\omega = (a^{-k}, \dots, a^{-1}) \in A^k$, $s \in S$ and $a \in A$, the transition function of this chain is

$$\pi(\omega', s' \mid \omega, s) := \sigma(a \mid \omega, s)p(s' \mid s),$$

where $\omega' = (a^{-(k-1)}, \dots, a^{-1}, a)$ is obtained by shifting the entries of ω , and appending a .

Therefore, a distribution $\nu \in \Delta(A^k \times S \times A)$ is invariant for (the Markov chain induced by) σ if for each $\omega = (a^{-k}, \dots, a^{-1}) \in A^k$ and $s \in S$,

$$\nu(\omega, s) = \sum_{a^{-(k+1)} \in A, s^{-1} \in S} \nu(\omega^{-1}, s^{-1})\sigma(a^{-1} \mid \omega^{-1}, s^{-1})p(s \mid s^{-1}),$$

where ω^{-1} stands for $(a^{-(k+1)}, \dots, a^{-2})$, and $\nu(\omega, s)$ is the marginal of ν over $A^k \times S$. Equivalently, a distribution $\nu \in \Delta(A^k \times S \times A)$ is invariant for *some* reporting policy σ if

$$\sum_a \nu(\omega, s, a) = \sum_{a^{-(k+1)} \in A, s^{-1} \in S} \nu(\omega^{-1}, s^{-1}, a^{-1})p(s \mid s^{-1}), \quad (12)$$

for all ω and s . Indeed, ν is then invariant for the policy σ defined by

$$\sigma(a \mid a^{-k}, \dots, a^{-1}, s) = \frac{\nu(a^{-k}, \dots, a^{-1}, s, a)}{\sum_{a' \in A} \nu(a^{-k}, \dots, a^{-1}, s, a')},$$

if the denominator is non-zero (and arbitrarily otherwise). We denote by \mathcal{N}_k the compact convex set of such distributions ν . We recall from Definition 3 that \mathcal{M}_k is the set of joint distributions $\mu \in \Delta(S \times A)$, for which there exists $\nu \in \mathcal{N}_k$, such that (i) the marginal of ν over $S \times A$ is μ and (ii) the marginal of ν over $A^k \times A$ is equal to the invariant distribution of (s_{n-k}, \dots, s_n) :

$$\sum_{s \in S} \nu(a^{-k}, \dots, a^{-1}, s, a) = \lambda(a^{-k})p(a^{-(k-1)} \mid a^{-k}) \cdots p(a \mid a^{-1}),$$

for all a^{-k}, \dots, a . Note that Equation (12) is the same as Equation (4) in the statement of Lemma 1. This completes the proof.

A.1.3 Proof of Lemma 2

For this proof, we will write $\mathcal{M}_k(p)$ to highlight the dependence on the transition matrix p . First consider the statement about dimension. Obviously $\mathcal{M}_k(p) \subset \mathcal{M}_0(p)$, which has dimension $(|S| - 1)^2$ due to the marginal constraints. To show this maximal dimension is achieved, we recall the linear system (2)–(5) defining \mathcal{M}_k . By Lemma 14 in Appendix D.1, for any $\mu \in \mathcal{M}_0(p)$ there is a (not necessarily positive) solution ν , which is zero whenever λ is.

Let $\mu^{ind}(s, a) = \lambda(s) \cdot \lambda(a)$ be the joint distribution achieved by reporting an identical Markov chain independent of the states. This policy induces a $\nu^{ind}(a^{-k}, \dots, a^{-1}, s, a)$ that is positive and strictly positive whenever $\lambda(a^{-k}, \dots, a^{-1}, a) > 0$. Thus if (ν, μ) solves (2)–(5), then for sufficiently small positive ε , $((1 - \varepsilon)\nu^{ind} + \varepsilon\nu, (1 - \varepsilon)\mu^{ind} + \varepsilon\mu)$ also satisfies the linear system. Since $(1 - \varepsilon)\nu^{ind} + \varepsilon\nu \geq 0$, this implies $(1 - \varepsilon)\mu^{ind} + \varepsilon\mu \in \mathcal{M}_k(p)$. Hence $\mathcal{M}_k(p)$ has the same dimension as $\mathcal{M}_0(p)$.

Next we prove continuity. U.h.c. is immediate because the constraints (2)–(5) vary continuously with p . To show l.h.c., we characterize $\mathcal{M}_k(p)$ by a dual linear program. Specifically, we can write (2)–(5) abstractly as

$$C(p) \cdot \nu = D(p, \mu), \nu \geq 0.$$

Here $C(p)$ is a finite $(c_1 \times c_2)$ matrix whose entries are linear combinations of entries in p ; $D(p, \mu)$ is a $c_1 \times 1$ column vector whose entries are either $\mu(s, a)$ or $\lambda(a^{-k})p(a^{-(k-1)}|a^{-k}) \dots p(a|a^{-1})$ or 0. To get rid of redundant constraints, let us only include those $\mu(s, a)$ where s, a belong to the first $|S| - 1$ states. In other words, we only consider $(|S| - 1)^2$ equations of type (2). Let $\tilde{\mu} \in \mathbf{R}^{(|S|-1)^2}$ be the projection of μ onto the first $|S| - 1$ states and reports. Similarly define $\tilde{\mathcal{M}}_k(p)$. Henceforth we write $D(p, \tilde{\mu})$ in place of $D(p, \mu)$.

By Farkas' lemma, there exists a non-negative solution ν to $C(p) \cdot \nu = D(p, \tilde{\mu})$ if and only if for any $y \in \mathbf{R}^{c_1}$,

$$y' \cdot C(p) \geq 0 \implies y' \cdot D(p, \tilde{\mu}) \geq 0. \quad (13)$$

Due to homogeneity, we can restrict attention to those y whose coordinates lie in $[-1, 1]$. Then the condition (13) simplifies to

$$y' \cdot D(p, \tilde{\mu}) \geq 0, \forall y' \in W(p), \quad (14)$$

where $W(p)$ is the bounded polytope $\{y \in \mathbf{R}^{c_1} : |y_j| \leq 1, y' \cdot C(p) \geq 0\}$, which is u.h.c. with respect to p . The above condition characterizes $\tilde{\mathcal{M}}_k(p)$ via a family of linear constraints on $\tilde{\mu}(s, a)$. In fact, only a finite collection of constraints matters, because $W(p)$ has finitely many vertices.

To prove l.h.c., we now fix p as well as a sequence $p_n \rightarrow p$. Fix $\tilde{\mu} \in \text{int}(\tilde{\mathcal{M}}_k(p))$, whose existence is guaranteed by dimension-counting. Then the finitely many relevant constraints in (14) cannot be binding at $\tilde{\mu}$. This implies the existence of $\varepsilon > 0$ such that

$$y' \cdot D(p, \tilde{\mu}) > \varepsilon, \forall y' \in W(p).$$

Take any $\tilde{\mu}_n \rightarrow \tilde{\mu}$ in $\mathbf{R}^{(|S|-1)^2}$. By the continuity of $D(p, \tilde{\mu})$ in both arguments, as well as the upper-hemicontinuity

of $W(p)$, we deduce that for all large n ,

$$y'_n \cdot D(p_n, \tilde{\mu}_n) > 0, \forall y'_n \in W(p_n).$$

By condition (14), this implies $\tilde{\mu}_n \in \text{int}(\tilde{\mathcal{M}}_k(p_n))$ for all large n . Thus the interior of $\tilde{\mathcal{M}}_k(p)$ satisfies l.h.c. The result follows because l.h.c. is preserved under closure.

A.2 Unprofitable deviations: proof of Proposition 2

We start with the direct implication. Assume truth-telling is k -limit optimal. Let $\mu \in \mathcal{M}_k$, and let (ν, σ) be a pair associated with μ . Since ν is invariant for σ , $\mu_{\nu, \sigma}^N = \mu$, hence $V^N(\nu, \sigma) = \mathbf{E}_\mu[r(s, a)]$ for each N . Since $(\nu, \sigma) \in \Sigma_k$ and by k -limit optimality, one has $\liminf_N V^N(\nu, \sigma) \leq \mathbf{E}_{\mu^{tt}}[r(s, a)]$, hence $\mathbf{E}_\mu[r(s, a)] \leq \mathbf{E}_{\mu^{tt}}[r(s, a)]$, as desired. Assume now that the inequality $\mathbf{E}_\mu[r(s, a)] \leq \mathbf{E}_{\mu^{tt}}[r(s, a)]$ holds for each $\mu \in \mathcal{M}_k$. Let $(q, \sigma) \in \Sigma_k$, and observe that $\liminf_{N \rightarrow \infty} V^N(q, \sigma) \leq \liminf_{\delta \rightarrow 1} V^\delta(q, \sigma)$. By Lemma 8, all limit points of $(\mu_{q, \sigma}^\delta)_\delta$ belong to \mathcal{M}_k . Since $V^\delta(q, \sigma) = \mathbf{E}_{\mu_{q, \sigma}^\delta}[r(s, a)]$ for each δ , this yields

$$\liminf_{\delta \rightarrow 1} V^\delta(q, \sigma) \leq \sup_{\mu \in \mathcal{M}_k} \mathbf{E}_\mu[r(s, a)],$$

and therefore, as desired,

$$\liminf_{N \rightarrow \infty} V^N(q, \sigma) \leq \mathbf{E}_{\mu^{tt}}[r(s, a)].$$

A.3 On the sufficiency of singleton states: proof of Theorem 1

The more difficult implication is the reverse one, which is proven in Renault, Solan, and Vieille (2013). They show that if $(s_n)_n$ is pseudo-renewal, then for any $\mu \in \mathcal{M}_0$ there is a reporting policy such that (i) the law of the entire sequence of reports $(a_n)_n$ is equal to the law of the sequence $(s_n)_n$ of states and (ii) the distribution of the state-report pair (s_n, a_n) is equal to μ for each n . It follows that $\mu \in \mathcal{M}_k$, so that $\mathcal{M}_0 = \mathcal{M}_k$ for each k .

Assume now that the sequence $(s_n)_n$ is not pseudo-renewal. For concreteness, assume that $S = \{s_1, \dots, s_{|S|}\}$. Up to a relabelling of the states, we may assume that $p(s_3 | s_1) \neq p(s_3 | s_2)$. Consider then the permutation matrix Π that permutes the two states s_1 and s_2 . For $\varepsilon \leq \min(\lambda_1, \lambda_2)$, the distribution

$$\mu_\varepsilon := \lambda + \varepsilon(\Pi - I) = \begin{pmatrix} \lambda_1 - \varepsilon & \varepsilon & & & \\ \varepsilon & \lambda_2 - \varepsilon & & & \\ & & \lambda_3 & & \\ & & & \ddots & \\ & & & & \lambda_{|S|} \end{pmatrix}$$

(where λ is treated as a diagonal matrix) belongs to \mathcal{M}_0 .

We claim that, for $\varepsilon > 0$, μ_ε fails to belong to \mathcal{M}_1 . Else, there would exist a reporting policy $\sigma : A \times S \rightarrow \Delta(A)$, and an invariant distribution $\nu \in \Delta(A \times S \times A)$, such that properties (i) and (ii) in Definition 3 hold.

By (ii), the frequency of rounds in which the agent reports successively (a_1, a_3) is $\nu(a_1, a_3) = \lambda_1 p(s_3 | s_1)$. By (i), and since $a = a_3$ if and only if $s = s_3$, it is also equal to

$$\begin{aligned} \nu(a_1, a_3) &= \mathbf{P}_\sigma(s_1, a_1, a_3) + \mathbf{P}_\sigma(s_2, a_1, a_3) \\ &= (\lambda_1 - \varepsilon)p(s_3 | s_1) + \varepsilon p(s_3 | s_2) \\ &\neq \lambda_1 p(s_3 | s_1). \end{aligned}$$

Thus, the half-line starting at μ^{tt} and pointing in the direction $\Pi - I$ only intersects \mathcal{M}_1 at μ^{tt} , but does intersect \mathcal{M}_0 along a non-trivial segment. In other words, $\mathcal{C}_0 \not\supseteq \mathcal{C}_1$, as we need to show.

B Implications and Applications

B.1 Proof of Proposition 7

We start with the reverse implication. Assume that truth-telling is optimal for some stationary transfer $t : A^{k+1} \rightarrow \mathbf{R}$ with memory k . Let an arbitrary $\mu \in \mathcal{M}_k$ be given, let σ be a reporting policy, and $\nu \in \Delta(A^k \times S \times A)$ be an invariant measure for σ associated with μ . Because truth-telling is optimal,

$$\mathbf{E}_\nu [r(s, \phi(a)) + t(a^{-k}, \dots, a^{-1}, a)] \leq \mathbf{E}_{\mu^{tt}} [r(s, \phi(a)) + t(a^{-k}, \dots, a^{-1}, a)].$$

Because $\mu \in \mathcal{M}_k$, expected transfers are the same on both sides, and the latter inequality rewrites, as desired,

$$\mathbf{E}_\mu [r(s, \phi(a))] \leq \mathbf{E}_\lambda [r(s, \phi(s))].$$

The proof of the direct implication relies on a minmax theorem. Assume that truth-telling is k -limit optimal, so that $\mathbf{E}_\mu [r(s, \phi(a))] \leq \mathbf{E}_\lambda [r(s, \phi(s))]$ for each $\mu \in \mathcal{M}_k$. Consider the zero-sum game in which the designer chooses a transfer $t : A^{k+1} \rightarrow \mathbf{R}$ such that $\mathbf{E}_{\mu^{tt}} [t(a^{-k}, \dots, a^{-1}, a)] = 0$, the agent chooses an invariant distribution $\nu \in \mathcal{N}_k$ that arises under some stationary policy σ , and the payoff to the agent is

$$g(\nu, t) := \mathbf{E}_\nu [r(s, \phi(a)) + t(a^{-k}, \dots, a^{-1}, a)].$$

Both pure strategy sets are convex, and the agent's strategy set is compact. Since the payoff function is affine in each strategy, the game has a value in pure strategies. So, the agent has an optimal pure policy, by Sion Theorem.

Next, we claim that the value $V = \max_\nu \inf_t g(\nu, t)$ of the game is equal to $\mathbf{E}_\lambda [r(s, \phi(s))]$. Plainly, the agent can guarantee this amount by reporting truthfully, hence $V \geq \mathbf{E}_\lambda [r(s, \phi(s))]$. Fix now $\nu \in \mathcal{N}_k$. Assume first that the marginal over A^{k+1} of ν coincides with λ . In that case, the marginal μ of ν over $S \times A$ belongs to \mathcal{M}_k , and, for each t , one has

$$g(\nu, t) = \mathbf{E}_\mu [r(s, \phi(a))] \leq \mathbf{E}_\lambda [r(s, \phi(s))].$$

Assume now that the marginal of ν over A^{k+1} is not equal to λ . By a separation argument, there exists $t : A^{k+1} \rightarrow$

\mathbf{R} such that $\mathbf{E}_{\mu^{tt}}[t(a^{-k}, \dots, a^{-1}, a)] = 0$ and $\mathbf{E}_{\nu}[t(a^{-k}, \dots, a^{-1}, a)] < 0$. In that case, $\lim_{c \rightarrow +\infty} g(\nu, ct) = -\infty$. This concludes the proof that $V = \mathbf{E}_{\lambda}[r(s, \phi(s))]$.

We next show that the designer has an optimal policy. Given $\mu \in \Delta(S \times A)$, we denote by $d(\mu, \mathcal{M}_k)$ its (Euclidean) distance to the convex set \mathcal{M}_k . Given $\nu \in \mathcal{N}_k$, we denote by $\mu(\nu)$ and $\lambda(\nu)$ its marginals over $S \times A$ and A^{k+1} respectively. We omit the proof of the next claim. It relies on \mathcal{M}_k and \mathcal{N}_k being polytopes, and on the fact that $\mu(\nu) \in \mathcal{M}_k$ if and only if $\lambda(\nu) = \lambda$.

Claim 1 *There exists $c > 0$ such that $d(\mu(\nu), \mathcal{M}_k) \leq c\|\lambda(\nu) - \lambda\|_2$ for every $\nu \in \mathcal{N}_k$.*

This claim implies that for each ν , there is a pure policy t of the designer with $\|t\|_2 \leq c\|r\|_2$, such that $g(\nu, t) \leq V$. This is indeed obvious if $\lambda(\nu) = \lambda$. If not, set

$$t := -c \frac{\|r\|_2}{\|\lambda(\nu) - \lambda\|_2} (\lambda(\nu) - \lambda).$$

Denoting by $\mu \in \mathcal{M}_k$ the projection of $\mu(\nu)$ onto \mathcal{M}_k , one then has

$$\begin{aligned} r \cdot (\mu(\nu) - \lambda) &\leq \|r\|_2 \|\mu(\nu) - \lambda\|_2 \\ &\leq c\|r\|_2 \|\lambda(\nu) - \lambda\|_2 = \mathbf{E}_{\lambda}(t) - \mathbf{E}_{\nu}(t), \end{aligned}$$

where the first inequality uses the Cauchy-Schwarz inequality, and the second one follows from the previous claim.

Reorganizing terms, we get, as desired,

$$\mathbf{E}_{\nu}[r(s, \phi(a)) + t(a^{-k}, \dots, a)] \leq V.$$

This implies in turn that the zero-sum game in which the designer is restricted to the compact set of pure strategies t such that $\|t\|_2 \leq c\|r\|_2$ still has value V in pure strategies. In this restricted game, the designer has an optimal strategy, t . This concludes the proof of Proposition 7.

B.2 General mechanisms

Here, we provide some background details on Lemma 6 and 7.

A mechanism is a triple (M, T, Φ) , where $M = (M_n)_{n \geq 1}$ is the set of possible reports in each round (set $M_n = S$ for $n \leq 0$ in the case of a doubly-infinite horizon, as in Section 4.1), $T = (t_n)_n$ is the sequence of transfer functions, and $\Phi = (\Phi_n)_n$ is the sequence of allocation functions. For $n \geq 1$, we let \overleftarrow{M}_{n-1} stand for $\prod_{-\infty < k \leq n-1} M_k$: a generic element of \overleftarrow{M}_{n-1} is a sequence of reports up to round n , so that t_n and Φ_n map $\overleftarrow{M}_{n-1} \times M_n$ into \mathbf{R} and Y respectively. We assume that M_n is finite for each n , and endow \overleftarrow{M}_n with the product σ -field. To make sure expectations are well-defined, we impose that t_n and Φ_n be measurable for each n , and require in addition that (t_n) be bounded. Fixing a payoff criterion (either discounting or limit of means), a given mechanism (M, T, Φ) defines an optimization problem for the agent, denoted $\mathcal{P}_{M, T, \Phi}$. This problem is a (non-stationary) Markov decision problem, with state space $\overleftarrow{M}_{n-1} \times S$ and action set M_n in round $n \geq 1$; the

overall state space is thus defined to be the disjoint union of the sets $\overleftarrow{M}_{n-1} \times S$ for $n \geq 1$. Transitions and rewards in each round n are defined in the obvious way. An initial state in $\mathcal{P}_{M,T,\Phi}$ is the infinite sequence $\overleftarrow{s}_1 = (s_n)_{n \leq 1}$ of states up to and including round one.

A pure reporting policy (or, more simply, a policy) is a sequence $(\sigma_n)_{n \geq 1}$, where σ_n is a (measurable) map from $\overleftarrow{M}_{n-1} \times S$ into M_n that specifies which report to send in round n . Given a pure policy, the sequence of messages can be reconstructed unambiguously from the sequence $(\overleftarrow{s}_n)_n$ of states. Abusing notation, we may write either $m_n = \sigma_n(\overleftarrow{m}_{n-1}, s_n)$, or $m_n = \sigma_n(\overleftarrow{s}_n)$. Given a distribution q over initial states, we denote by \mathbf{P}_q the law of the entire sequence (s_n) of states, and by $\mathbf{P}_{q,\sigma}$ the law of the sequence $(s_n, m_n)_n$ under the policy σ . A map $\phi : S \rightarrow Y$ is *implementable* (or simply IC) if there exists a mechanism (M, T, Φ) and an optimal pure policy σ in $\mathcal{P}_{M,T,\Phi}$ such that $\phi(s_n) = \Phi_n(\sigma_n(\overleftarrow{s}_n))$, \mathbf{P} -a.s. for each n . A mechanism (M, T, Φ) is *direct* if $M_n = S$ for each $n \geq 1$ and if truth-telling $(\sigma_n(\overleftarrow{s}_n) = s_n$ for each $n)$ is an optimal policy in $\mathcal{P}_{M,T,\Phi}$.³⁰ Independent of the payoff criterion, the following revelation principle is immediate.

A mechanism (M, T, Φ) has *memory* $k \in \mathbf{N}$ if all maps t_n, Φ_n only depend on the current report m_n and on the k most recent ones, m_{n-k}, \dots, m_{n-1} . A policy σ in $\mathcal{P}_{M,T,\Phi}$ has memory k if, for each n , σ_n is a map from $M_{n-k} \times \dots \times M_{n-1} \times S$ into M_n . A mechanism, or policy, has *finite* memory if it has memory k for some $k < \infty$.

If $\phi : S \rightarrow Y$ is implementable via a mechanism with memory k , it need not be implementable via a *direct* mechanism with memory k . In other words, the proof of Lemma 6, when applied to a mechanism with memory k , does not deliver a direct mechanism with memory k , unless $k = 0$. The reason is that, given a mechanism (M, T, Φ) and a policy σ in $\mathcal{P}_{M,T,\Phi}$, both with memory k , σ_n need not have finite memory when viewed as a function $\sigma_n(\overleftarrow{s}_n)$ of previous states. As an example, set $S = \{0, 1\}$, $M_n = S$ for each n , and consider the reporting policy defined by $\sigma_1(s_1) = 0$ and $\sigma_n(m_{n-1}, s_n) = s_n + m_{n-1} \bmod 2$. While σ has memory 1 when viewed as a function of previous messages, it has unbounded memory when viewed as a function of previous states.

B.3 Proof of Lemma 5

Consider the set E of all mechanisms (M, T, Φ) with memory k such that $M_n = S$ and $\Phi_n(m_{n-k}, \dots, m_n) = \phi(m_n)$. The set E includes all such mechanisms which implement ϕ , and possibly others, as we do not require truth-telling to be optimal. The set E can be identified with the set of all sequences T of transfer functions with memory k , that is, with the set of all bounded sequences $T = (t_n)_{n \geq 1}$ in $\mathbf{R}^{S^{k+1}}$. Let $\Delta \subset E$ denote the subset of *stationary* mechanisms. We will prove that Δ contains a direct mechanism.

Let $a \in (0, 1)$ be arbitrary, and let η be the geometric distribution over \mathbf{N} with parameter a . We denote by $l^1(\eta)$ the normed space of sequences (t_n) in $\mathbf{R}^{S^{k+1}}$ such that

$$\mathbf{E}_\eta \|t_n\| = (1 - a) \sum_{n=1}^{+\infty} a^{n-1} \|t_n\| < +\infty.$$

Both E and Δ are linear subspaces of $l^1(\eta)$. In addition, Δ is a closed set (but E isn't).

Let $T^* = (t_n^*)_{n \geq 1} \in E$ be a *direct* mechanism and set $K := \sup_{n \geq 1} \|t_n^*\| < +\infty$. By time-invariance of the

³⁰Standard arguments show that, independent of the payoff criterion, optimal pure policies always exist.

truth-telling incentives, each of the shifted sequences $T^{*,l} = (t_{n+l}^*)_{n \geq 1}$ ($l \geq 0$) is also a direct mechanism. We denote by C the closure in $l^1(\eta)$ of the convex hull of the set $\{T^{*,l}, l \geq 0\}$.

Claim 2 *The set C is a compact subset of $l^1(\eta)$. All elements of C are direct mechanisms.*

Proof. Let an arbitrary element $T = (t_n)_n$ in C be given. Let $(T_p)_p$ be a sequence which converges to T (in the $l^1(\eta)$ sense), and such that each T_p is a convex combination of finitely many elements of the set $\{T^{*,l}, l \geq 0\}$. Writing $T_p = (t_{p,n})_n$, it follows from $\lim_{p \rightarrow \infty} T_p = T$ that $t_n = \lim_{p \rightarrow \infty} t_{p,n}$ for each $n \geq 1$. In particular, $\|t_n\| \leq K$ for each n .

In addition, and since truth-telling incentives are preserved under convex combinations of transfers, T_p is a direct mechanism for each p . By the dominated convergence theorem, the truth-telling incentives still hold at the limit $p \rightarrow +\infty$, and T is a direct mechanism.

Let $B \subset E$ be the set of sequences $T = (t_n)$, such that $\|t_n\| \leq K$ for all n . It follows from the first paragraph that $C \subset B$. By Tychonov Theorem, B is compact in the product topology. Since the topology induced by $l^1(\eta)$ on B coincides with the product topology on B , B is a compact subset of $l^1(\eta)$. Since C is a closed subset of $l^1(\eta)$ contained in B , it is compact as well. ■

Claim 3 *One has $C \cap \Delta \neq \emptyset$.*

Proof. Assume for the sake of contradiction that $C \cap \Delta = \emptyset$. Since C is compact and Δ is closed in $l^1(\eta)$, by the hyperplane separation theorem, there exist two real numbers $c_1 > c_2$, and a continuous linear form p on $l^1(\eta)$ such that $p(T) \geq c_1$ and $p(T') \leq c_2$ for every $T \in \Delta$ and $T' \in C$. By Riesz Theorem, p may be identified with a bounded sequence (p_n) in $\mathbf{R}^{S^{k+1}}$, with $p(T) = (1-a) \sum_{n=1}^{+\infty} a^{n-1} p_n \cdot t_n$ for each $T \in l^1(\eta)$, where \cdot denotes the standard scalar product on $\mathbf{R}^{S^{k+1}}$.

Since Δ is a linear subspace, the condition $p(T) \geq c_1$ for each $T \in \Delta$ implies $c_1 = 0$ and $p(T) = 0$ for each $T \in \Delta$. In particular, one has

$$(1-a) \sum_{n=1}^{+\infty} a^{n-1} p_n^i = 0, \quad (15)$$

for each component $i \in S^{k+1}$. On the other hand, one has $p(T^l) \leq c_2 < 0$ for each $l \in \mathbf{N}$, which implies

$$p(T^0) + \cdots + p(T^l) \leq (l+1)c_2.$$

We conclude by proving that the left-hand side $p(T^0) + \cdots + p(T^l)$ is bounded, which yields the desired contradiction since $(l+1)c_2 \rightarrow -\infty$. Set $K' = \sup_{n \geq 1} \|p_n\|$, so that $|p_n \cdot t_m| \leq K_1 := |S^{k+1}| K K'$ for each n and m . One has

$$\begin{aligned} \sum_{j=0}^l p(T^j) &= (1-a) \sum_{j=0}^l \sum_{n=1}^{+\infty} a^{n-1} p_n \cdot t_{n+j} \\ &= (1-a) \sum_{m=1}^{+\infty} \sum_{n=\max\{1, m-l\}}^m a^{n-1} p_n \cdot t_m \\ &= (1-a) \sum_{m=1}^l \sum_{n=1}^m a^{n-1} p_n \cdot t_m + (1-a) \sum_{m=l+1}^{+\infty} \sum_{n=m-l}^m a^{n-1} p_n \cdot t_m. \end{aligned} \quad (16)$$

By (15), for each m one has $\sum_{n=1}^m a^{n-1} p_n \cdot t_m = - \sum_{n=m+1}^{+\infty} a^{n-1} p_n \cdot t_m$. Therefore,

$$\left| (1-a) \sum_{m=1}^l \sum_{n=1}^m a^{n-1} p_n \cdot t_m \right| \leq (1-a) \sum_{m=1}^l a^m K_1 \leq a K_1.$$

As for the second term on the R.H.S. of (16), one has

$$\left| (1-a) \sum_{m=l+1}^{+\infty} \sum_{n=m-l}^m a^{n-1} p_n \cdot t_m \right| \leq (1-a) \sum_{m=l+1}^{+\infty} \sum_{n=m-l}^m a^{n-1} K_1 = K_1 \frac{1-a^{l+1}}{1-a}.$$

Both of these terms are bounded, which concludes the proof. ■

B.4 Proofs of Lemmata 6 and 7

The proof of Lemma 6 follows standard arguments. Fix a mechanism (M, T, Φ) and an optimal pure policy σ in $\mathcal{P}_{M, T, \Phi}$ that implements ϕ . For $n \geq 1$, define $t'_n(\overleftarrow{s}_n) = t_n(\overleftarrow{\sigma_n(\overleftarrow{s}_n)})$. From the optimality of σ in $\mathcal{P}_{M, T, \Phi}$, it follows that truth-telling is optimal in $\mathcal{P}_{A, T', \phi}$.³¹

We turn to the proof of Lemma 7. Let (M, T, Φ) be a given mechanism with memory k , with a pure optimal policy σ that implements ϕ . We construct a memory 1 mechanism that implements ϕ . Intuitively, in addition to reporting m_n , the agent has to repeat in round n the reports made in the k previous rounds. The reports in two consecutive rounds relative to any given, earlier round will be incentivized to coincide.

To be formal, we set $M'_n = M_{n-k} \times \dots \times M_n$, $\Phi'_n(m'_{n-1}, m'_n) = \Phi_n(m'_n)$ for each $n \geq 1$, and we define $t'_n(m'_{n-1}, m'_n)$ to be equal to $t_n(m'_n)$ if the sequences m'_{n-1} and m'_n of reports are consistent, and to be a large negative number P otherwise.³²

If the penalty P is severe enough, optimal policies in $\mathcal{P}_{M', T', \Phi'}$ send consistent messages in any two consecutive rounds. When restricted to such “consistent” policies, the incentives faced by the agent in $\mathcal{P}_{M, T, \Phi}$ and in $\mathcal{P}_{M', T', \Phi'}$ are identical. Therefore, the policy σ' , with $\sigma'_n(m'_{n-1}, s_n) = (m_{n-k}, \dots, m_{n-1}, \sigma_n(m_{n-k}, \dots, m_{n-1}, s_n))$ is an optimal policy in $\mathcal{P}_{M', T', \Phi'}$ that implements ϕ .

References

- [1] Altman, E. (1999). *Constrained Markov Decision Processes*, Chapman & Hall/CRC.

³¹Slightly abusing notations: in $\mathcal{P}_{A, T', \phi}$, the message space is A in each round, and $\phi_n = \phi$ for each n .

³²Here, $m'_{n-1} = (m_{n-k-1}, \dots, m_{n-1})$.

Keeping Your Story Straight: Truth-telling and Liespotting

J. Hörner, X. Mu, and N. Vieille

April, 2017

Supplementary Online Appendix

C Discussion of the Assumptions

C.1 Discounting

We here prove Proposition 3 by following the proof of Proposition 7. Note that in the zero-sum game considered in that proof, any optimal policy ν_* of the agent is such that $\lambda(\nu_*) = \lambda$. In particular, under the stronger assumption $\mathbf{E}_\mu[r(s, \phi(a))] < \mathbf{E}_\lambda[r(s, \phi(s))]$ for each $\mu \in \mathcal{M}_k \setminus \{\lambda\}$, truth-telling is the *unique* optimal policy of the agent in this zero-sum game.

Recall that $\nu_{q,\sigma}^\delta \in \Delta(A^k \times S \times A)$ is the expected discounted frequency given (q, σ) . Let us set

$$\mathcal{N}_k^\delta = \text{co}\{\nu_{q,\sigma}^\delta, \text{ all } q, \sigma\}.$$

Lemma 10 *As $\delta \rightarrow 1$, \mathcal{N}_k^δ converges to \mathcal{N}_k (see the proof of Lemma 1) in the Hausdorff distance.*

Proof. Let $\nu \in \mathcal{N}_k$ be given, and σ be a stationary reporting policy which admits ν as an invariant measure. For each δ , if the agent follows σ and the initial distribution is ν , then expected discounted frequencies coincide with ν . Hence \mathcal{N}_k is a subset of $\lim_{\delta \rightarrow 1} \mathcal{N}_k^\delta$.

Conversely, it is sufficient to prove that for any $(\delta_m)_m \rightarrow 1$ and $(\nu_m)_m \rightarrow \nu$, with $\nu_m \in \mathcal{N}_k^{\delta_m}$ for each m , one has $\nu \in \mathcal{N}_k$. This follows from Lemma 9. ■

Proposition 3 follows easily from Lemma 10 and the proof of Proposition 7. Given $\delta < 1$, consider the zero-sum game in which the action set of the agent is \mathcal{N}_k^δ , the action set of the principal is the set of t such that $\|t\|_2 \leq c$ and $\mathbf{E}_\lambda[t(s^{-k}, \dots, s)] = 0$, and the payoff is defined as in the proof of Proposition 7. The game has a value V^δ , and we denote by ν_*^δ an arbitrary optimal policy of the agent. It follows from the Hausdorff convergence of \mathcal{N}_k^δ to \mathcal{N}_k that V^δ converges to V , and that any limit point of $(\nu_*^\delta)_\delta$ as $\delta \rightarrow 1$ is an optimal policy of the agent in the zero-sum game of Proposition 2. As stressed above, truth-telling is the unique optimal policy in that game. Therefore, the marginal μ^δ of ν_*^δ converges to μ^{tt} .

C.2 Non-Anticipation

Lemma 7 *Suppose that the agent is $|S| + 1$ -prophetic. Then truth-telling can be achieved if, and only if, it can be achieved when checking singleton states.*

Proof. We write the proof under the additional assumption that $p(\cdot | s)$ has full support for each s . Assume that truth-telling is optimal for some stationary transfer t with memory k . Let $\vec{s} = (s^1, \dots, s^p)$, $p \leq |S|$ be any list of p distinct states. We identify the list \vec{s} with a permutation π over S , which leaves fixed all remaining states. Denote by Π the permutation matrix associated with π .

Under the assumption that the agent is $|S| + 1$ -prophetic, we will construct below a non-stationary reporting policy σ , by blocks of size $p + 1$, and with the following properties:

P1 the sequence $(a_n)_n$ of reports has the same probability distribution as the sequence $(s_n)_n$ of states,

P2 the joint distribution of the state-report pair (s_n, a_n) in a given round n is equal to λ if n is a multiple of $p + 1$, and is equal to $\lambda + \varepsilon(\Pi - I)$ otherwise (for some positive ε).

Since the sequences $(a_n)_n$ and $(s_n)_n$ have the same distribution, the expected transfer to the agent in any round n is the same, whether he reports truthfully or according to σ . Since truth-telling is optimal, this implies that the expectation of $r(s, \phi(a))$ under Π does not exceed the expectation under I , that is,

$$\sum_{j=1}^p r(s^j, \phi(s^{j+1})) \leq \sum_{j=1}^p r(s^j, \phi(s^j)).$$

Since this is true for each list \vec{s} , this will imply that ϕ is 0-IC.

For $j = 0, \dots, p$, we denote by \vec{s}^j the list $(s^j, \dots, s^p, s^1, \dots, s^{j-1})$ obtained from \vec{s} , when starting from its j -th element. Pick $\varepsilon > 0$ small enough so that $\mathbf{P}(s, \vec{s}^j, s') \geq \varepsilon$ for all $j, s, s' \in S$ —that is, the long-run frequency of the sequence (s, \vec{s}^j, s') of length $p + 2$ is at least ε . If the sequence of forthcoming states s_{n+1}, \dots, s_{n+p} is equal to \vec{s}^j for some j , the agent reports (s_n, \vec{s}^{j+1}) with a probability $\frac{\varepsilon}{\mathbf{P}(\vec{s}^j | s_{iT}, s_{(i+1)T})}$, and reports truthfully (s_n, \vec{s}^j) with the residual probability. If the sequence of forthcoming states fails to coincide with any permutation \vec{s}^j of \vec{s} , the agent reports truthfully on this block.

Conditional on the first states s_n and s_{n+p+1} of two consecutive blocks, the probability of reporting \vec{s}^{j+1} rather than \vec{s}^j is equal to ε . Hence, conditional on s_n and s_{n+p+1} , the sequence of reports within the block has the same distribution as the distribution of states, and is independent of the reports submitted in the earlier blocks. This implies **P1**.

By construction, σ reports truthfully in the first round of each block. For any other round n , one has either $s_n = a_n$ or $(s_n, a_n) = (s^j, s^{j+1})$, for some j . The latter occurs with probability ε , for each j . This establishes **P2**. ■

D Properties of \mathcal{M}_k and Dual Programs

D.1 Iterative linear system

We recall the linear system defining \mathcal{M}_k :

$$\sum_{a^{-k}, \dots, a^{-1}} \nu(a^{-k}, \dots, a^{-1}, s, a) = \mu(s, a). \quad (2)$$

$$\sum_s \nu(a^{-k}, \dots, a^{-1}, s, a) = \lambda(a^{-k})p(a^{-(k-1)}|a^{-k}) \cdots p(a|a^{-1}). \quad (3)$$

$$\sum_a \nu(a^{-k}, \dots, a^{-1}, s, a) = \sum_{s^{-1}} \left(p(s|s^{-1}) \cdot \sum_{a^{-(k+1)}} \nu(a^{-(k+1)}, \dots, a^{-2}, s^{-1}, a^{-1}) \right). \quad (4)$$

$$\nu(a^{-k}, \dots, a^{-1}, s, a) \geq 0. \quad (5)$$

Given some $\mu \in \mathcal{M}_0$, we can determine whether $\mu \in \mathcal{M}_k$ by directly solving the above linear program. Alternatively, we can also check iteratively whether some undetectable distribution $\nu(a^{-(j-1)}, \dots, a^{-1}, s, a)$ under memory $j-1$ might be extended to some $\nu(a^{-j}, \dots, a^{-1}, s, a)$ that is undetectable under memory j . This is formally expressed in the next lemma:

Lemma 13 *$\mu \in \mathcal{M}_k$ if and only if there exists $\nu : \cup_{j=0}^k (A^j \times S \times A) \rightarrow \mathbf{R}_+$, such that $\nu(s, a) = \mu(s, a)$, and for $1 \leq j \leq k$ we have the following:*

$$\sum_{a^{-j}} \nu(a^{-j}, \dots, a^{-1}, s, a) = \nu(a^{-(j-1)}, \dots, a^{-1}, s, a). \quad (17)$$

$$\sum_s \nu(a^{-j}, \dots, a^{-1}, s, a) = \lambda(a^{-j})p(a^{-(j-1)}|a^{-j}) \cdots p(a|a^{-1}). \quad (18)$$

$$\sum_a \nu(a^{-j}, \dots, a^{-1}, s, a) = \sum_{s^{-1}} p(s|s^{-1}) \cdot \nu(a^{-j}, \dots, a^{-2}, s^{-1}, a^{-1}). \quad (19)$$

Proof. Given $\nu(a^{-k}, \dots, a^{-1}, s, a)$ that satisfies equations (2) to (5), simply define

$$\nu(a^{-j}, \dots, a^{-1}, s, a) = \sum_{a^{-k}, \dots, a^{-(j+1)}} \nu(a^{-k}, \dots, a^{-1}, s, a).$$

Conditions (17) to (19) are checked in a straightforward manner. ■

Although the proof is simple, this result is very useful in practice because the linear system (17) to (19) allows us to solve for $\nu(a^{-j}, \dots, a^{-1}, s, a)$ on the L.H.S. in terms of $\nu(a^{-(j-1)}, \dots, a^{-1}, s, a)$ on the R.H.S., a procedure that can be iterated. Importantly, note that we can hold $a^{-(j-1)}, \dots, a^{-1}$ as fixed and consider (17) to (19) as a *3-dimensional transportation problem* in the variables a^{-j}, s and a , see Smith and Dawson (1979). This way, we reduce the original large linear system (with $|S|^{k+2}$ variables) into a collection of smaller linear systems (each with $|S|^3$ variables). The following is an easy corollary:

Lemma 14 *For any $\mu \in \mathcal{M}_0$, there exists $\nu(a^{-k}, \dots, a^{-1}, s, a)$ (not necessarily positive) satisfying equations (2) to (4), and $\nu(a^{-k}, \dots, a^{-1}, s, a) = 0$ whenever $\lambda(a^{-k}, \dots, a^{-1}, a) = 0$.*

Proof. Let us first ignore the requirement that ν be zero when the corresponding λ is. By the previous lemma, it suffices to show that for fixed $\nu(a^{-(j-1)}, \dots, a^{-1}, s, a)$, we can solve for $\nu(a^{-j}, \dots, a^{-1}, s, a)$ (not necessarily positive) from equations (17) to (19). Holding fixed $a^{-(j-1)}, \dots, a^{-1}$, we can treat the linear system as a 3-dimensional transportation problem. A solution exists if and only if certain “add-up constraints” regarding the R.H.S. of (17) to (19) are satisfied. Specifically we need:

$$\begin{aligned} \sum_s \nu(a^{-(j-1)}, \dots, a^{-1}, s, a) &= \sum_{a^{-j}} \lambda(a^{-j}) p(a^{-(j-1)}|a^{-j}) \cdots p(a|a^{-1}). \\ \sum_a \nu(a^{-(j-1)}, \dots, a^{-1}, s, a) &= \sum_{a^{-j}} \sum_{s^{-1}} p(s|s^{-1}) \cdot \nu(a^{-j}, \dots, a^{-2}, s^{-1}, a^{-1}). \\ \sum_a \lambda(a^{-j}) p(a^{-(j-1)}|a^{-j}) \cdots p(a|a^{-1}) &= \sum_s \sum_{s^{-1}} p(s|s^{-1}) \cdot \nu(a^{-j}, \dots, a^{-2}, s^{-1}, a^{-1}). \end{aligned}$$

These follow from the fact that $\nu(a^{-(j-1)}, \dots, a^{-1}, s, a)$ is a solution to the linear system defining \mathcal{M}_{j-1} , which can be shown by induction.

To incorporate the additional requirement, we need to show that if $\nu(a^{-(j-1)}, \dots, a^{-1}, s, a) = 0$ whenever $\lambda(a^{-(j-1)}, \dots, a^{-1}, s, a) = 0$, then for fixed $a^{-(j-1)}, \dots, a^{-1}$ there exists a solution $\nu(a^{-j}, \dots, a^{-1}, s, a)$ to the transportation problem that is zero whenever λ is. If $\lambda(a^{-(j-1)}, \dots, a^{-1}) = 0$, we set $\nu(a^{-j}, \dots, a^{-1}, s, a)$ to be zero for all a^{-j}, s, a . Otherwise we only need $\nu(a^{-j}, \dots, a^{-1}, s, a)$ to be zero when $a^{-j} \in A_1$ or $a \in A_2$, for some $A_1, A_2 \subset A$. This way, we can ignore the value of ν at these points and treat equations (17) to (19) as a $|A_1^c| \times |S| \times |A_2^c|$ transportation problem. The add-up constraints are still satisfied, so we can find such a ν . ■

D.2 Dual constraints and direct characterization of \mathcal{M}_1 when $|S| = 3$

In proving the continuity of \mathcal{M}_k with respect to p , we considered the linear program dual to (2)–(5). Here we fully develop the dual program to obtain a *direct* characterization of \mathcal{M}_1 without any reference to $\nu(a^{-1}, s, a)$. We will focus on the case with 3 states, providing foundation for the proof of our main sufficiency/insufficiency results (Proposition ?? to 5).

We begin by recalling that $\mu \in \mathcal{M}_1$ if and only if there exists $\nu(a^{-1}, s, a) \geq 0$ satisfying:

$$\begin{aligned} \sum_{a^{-1}} \nu(a^{-1}, s, a) &= \mu(s, a). \\ \sum_s \nu(a^{-1}, s, a) &= \lambda(a^{-1})p(a|a^{-1}). \\ \sum_a \nu(a^{-1}, s, a) &= \sum_{s^{-1}} p(s|s^{-1})\mu(s^{-1}, a^{-1}). \end{aligned}$$

Given μ , this is a 3-dimensional transportation problem. By Farkas' lemma, the existence of a solution $\nu \geq 0$ is equivalent to μ having the correct marginals and satisfying a finite collection of linear inequality constraints, which we will call the “dual constraints”. To make the problem symmetric, let us define $f(s, a) = \mu(s, a)$, $g(a^{-1}, a) = \lambda(a^{-1})p(a|a^{-1})$, and $h(a^{-1}, s) = \sum_{s^{-1}} p(s|s^{-1})\mu(s^{-1}, a^{-1})$. When $|S| = 3$, we know from Smith and Dawson (1979) that the dual constraints are:

$$f, g, h \geq 0, \tag{20}$$

$$f(s, a) \leq g(a^{-1}, a) + h(\hat{a}^{-1}, s) + h(\tilde{a}^{-1}, s), A = \{a^{-1}, \hat{a}^{-1}, \tilde{a}^{-1}\}, \text{ similarly if we permute } f, g, h, \tag{21}$$

$$f(s, a) + f(s', a) \leq g(a^{-1}, a) + h(\hat{a}^{-1}, s) + h(\hat{a}^{-1}, s') + h(\tilde{a}^{-1}, s) + h(\tilde{a}^{-1}, s'), s' \neq s \text{ and permutations,} \tag{22}$$

$$f(s, a) + g(a^{-1}, a') + h(a^{-1}, s') \leq L(a^{-1}) + f(s', a') + g(\hat{a}^{-1}, a) + h(\tilde{a}^{-1}, s), s' \neq s, a' \neq a \text{ and permutations,} \quad (23)$$

where $L(a^{-1}) = \sum g(a^{-1}, \cdot) = \sum h(a^{-1}, \cdot)$.

To understand these constraints, we can visualize a $3 \times 3 \times 3$ cube filled with numbers $\nu(a^{-1}, s, a)$. Then $f/g/h$ are the sums along the $a^{-1}/s/a$ directions. (21) holds because the a^{-1} -direction “line” representing the sum on the L.H.S. is covered by the lines that appear on the R.H.S. As an example, consider $f(s_1, a_1) \leq g(a_1, a_1) + h(a_2, s_1) + h(a_3, s_1)$. If $\nu(a^{-1}, s, a)$ is a solution to the linear system, then the preceding inequality is equivalent to

$$\nu(a_1, s_2, a_1) + \nu(a_1, s_3, a_1) + \nu(a_2, s_1, a_2) + \nu(a_2, s_1, a_3) + \nu(a_3, s_1, a_2) + \nu(a_3, s_1, a_3) \geq 0.$$

Thus this particular constraint binds at some $\mu \in \mathcal{M}_1$ if and only if for this μ , the associated ν (which can be non-unique) has zeros in the following positions:

$$\nu = \begin{pmatrix} + & + & + \\ 0 & + & + \\ 0 & + & + \end{pmatrix} \begin{pmatrix} + & 0 & 0 \\ + & + & + \\ + & + & + \end{pmatrix} \begin{pmatrix} + & 0 & 0 \\ + & + & + \\ + & + & + \end{pmatrix}. \quad (24)$$

where the matrix represents a^{-1} , the row represents s and the column represents a . For ease of exposition, we often use figures of that type in the proofs that follow (“+” means we do not know whether it is zero or strictly positive). For future reference, we will call the constraint (21) a *type-I constraint*. Its distinguishing feature is that ν has to have zero in 6 positions.

Similarly, the following is an example where (22) binds:

$$\nu = \begin{pmatrix} 0 & + & + \\ + & + & + \\ + & + & + \end{pmatrix} \begin{pmatrix} + & + & + \\ + & 0 & 0 \\ + & 0 & 0 \end{pmatrix} \begin{pmatrix} + & + & + \\ + & 0 & 0 \\ + & 0 & 0 \end{pmatrix}. \quad (25)$$

We call this a *type-II constraint*, which requires ν to be zero in a $2 \times 2 \times 2$ sub-cube and also in one more position (that is not “co-planar” with the sub-cube).

Lastly we demonstrate a figure that shows (23) to be binding:

$$\nu = \begin{pmatrix} + & + & 0 \\ + & + & + \\ 0 & + & 0 \end{pmatrix} \begin{pmatrix} + & + & + \\ 0 & 0 & + \\ 0 & + & + \end{pmatrix} \begin{pmatrix} + & 0 & 0 \\ + & 0 & + \\ + & + & + \end{pmatrix}. \quad (26)$$

This is called a *type-III constraint*. Compared to a type-II constraint, type-III also requires ν to be zero in 9 positions. But the difference now is that on every 3×3 face of the cube, ν is zero in exactly 3 positions.

Finally we call the positivity constraint (20) a *type-0 constraint*. To summarize, we have shown the following:

Lemma 15 *When $|S| = 3$, \mathcal{M}_1 is the set of $\mu \in \mathcal{M}_0$ that satisfy the dual constraints (20) to (23), which are called type-0, I, II, III constraints. The latter three constraints bind if and only if the $\nu(a^{-1}, s, a)$ associated with μ has zeros in certain positions, as illustrated by (24) to (26).*

E Sufficiency of \mathcal{M}_1

In this appendix we prove the following result regarding the sufficiency of memory 1:

Proposition 3 *Let $|S| = 3$. Assume $p_{11} > p_{22} > p_{33} > \max\{p_{21}, p_{31}\} > \max\{p_{12}, p_{32}\} > \max\{p_{13}, p_{23}\} > 0$, and also $p_{11} + p_{22} < 1$. Then $\mathcal{M}_1 = \mathcal{M}_k$ for every $k \geq 2$.*

We will additionally assume $p_{21} > p_{31}, p_{12} > p_{32}$ and $p_{13} > p_{23}$, so that $p_{11} > p_{22} > p_{33} > p_{21} > p_{31} > p_{12} > p_{32} > p_{13} > p_{23} > 0$. Essentially the same proof strategy works for other cases as well, and we omit the details.

E.1 Proof outline and the polytopes \mathcal{M}_* , $\overline{\mathcal{M}_*}$

Before supplying too much detail, let us first sketch an outline of the proof. Our goal is to show that every vertex μ of \mathcal{M}_1 belongs to \mathcal{M}_k . For this we need to construct an undetectable reporting policy σ that achieves μ as a joint distribution. One such policy involves running a Markov chain on $S \times A$, such that the conditional probability of a given (s^{-1}, a^{-1}) is independent of s^{-1} and given by $p(a|a^{-1})$. If this holds, then an outsider could view (a_n) as autonomous, while s_n depends on s_{n-1}, a_{n-1} and a_n . Such a policy is undetectable for all memory k .

To formalize the above idea, we introduce the set \mathcal{M}_* :

Definition 7 \mathcal{M}_* is the set of joint distributions $\mu \in \Delta(S \times A)$ for which the following linear system in $\tau(s^{-1}, a^{-1}, s, a)$ has a solution:

$$\sum_{s^{-1}, a^{-1}} \tau(s^{-1}, a^{-1}, s, a) = \mu(s, a). \quad (27)$$

$$\sum_a \tau(s^{-1}, a^{-1}, s, a) = \mu(s^{-1}, a^{-1}) \cdot p(s|s^{-1}). \quad (28)$$

$$\sum_s \tau(s^{-1}, a^{-1}, s, a) = \mu(s^{-1}, a^{-1}) \cdot p(a|a^{-1}). \quad (29)$$

$$\tau(s^{-1}, a^{-1}, s, a) \geq 0.$$

To interpret, μ will be the invariant distribution of a Markov chain on $S \times A$, whose transition is given by $\frac{\tau(s^{-1}, a^{-1}, s, a)}{\mu(s^{-1}, a^{-1})}$. Equation (27) ensures that μ is indeed the invariant distribution. Equation (28) captures the assumption that (s_n) is autonomous. Equation (29) is the restriction that (a_n) could also be viewed as autonomous. When these conditions are satisfied, an agent reporting a with probability $\frac{\tau(s^{-1}, a^{-1}, s, a)}{p(s|s^{-1}) \cdot \mu(s^{-1}, a^{-1})}$ induces the desired Markov chain over $S \times A$, and he achieves the joint distribution μ in an undetectable way. We have thus shown $\mathcal{M}_* \subset \mathcal{M}_k$ for every k .

By this observation, we could prove $\mathcal{M}_1 = \mathcal{M}_k$ if it happens that $\mathcal{M}_1 = \mathcal{M}_*$. However, our computation suggests that the latter is only true for highly non-generic p .³³

To bridge the gap, we will next define a superset of \mathcal{M}_* that still lives in every \mathcal{M}_k . Given some τ that satisfies (27) to (29), we can define the induced distribution $\nu \in \Delta(A^k \times S \times A)$ iteratively as follows:

$$\nu(s, a) = \mu(s, a). \quad (30)$$

$$\nu(a^{-k}, \dots, a^{-1}, s, a) = \sum_{s^{-1}} \nu(a^{-k}, \dots, a^{-2}, s^{-1}, a^{-1}) \cdot \tau(s, a|s^{-1}, a^{-1}), \forall k \geq 1. \quad (31)$$

By the autonomous property of (s_n) and (a_n) , $\nu(a^{-k}, \dots, a^{-1}, s, a)$ satisfies the linear system in order for μ to be in \mathcal{M}_k , ignoring the positivity constraint. We thus only need

$$\nu(a^{-k}, \dots, a^{-1}, s, a) \geq 0. \quad (32)$$

to conclude that $\mu \in \mathcal{M}_k$. $\tau \geq 0$ is certainly a sufficient condition for the positivity of ν , but it

³³We can show that when $w \in [\frac{1}{4}, \frac{1}{2}]$ and all transition probabilities in p are either w or $1 - 2w$, then $\mathcal{M}_1 = \mathcal{M}_*$.

turns out to be unnecessary, motivating a relaxed definition:³⁴

Definition 8 $\overline{\mathcal{M}}_*$ is the set of joint distributions $\mu \in \Delta(S \times A)$ for which the system of equations (27) to (32) has a solution in τ and ν .

Lemma 16 $\mathcal{M}_* \subset \overline{\mathcal{M}}_* \subset \mathcal{M}_k$ for every k .

We omit the proof because it is already in the discussion. Using the above lemma, we can deduce Proposition ?? from the following alternative statement:

Proposition 3' Assume $p_{11} > p_{22} > p_{33} > p_{21} > p_{31} > p_{12} > p_{32} > p_{13} > p_{23} > 0$ and $p_{11} + p_{22} < 1$, the polytope \mathcal{M}_1 has 22 vertices. Each of these vertices belongs to $\overline{\mathcal{M}}_*$.³⁵

We devote future subsections to the proof of Proposition 3'.

E.2 Binding dual constraints

When there are 3 states, \mathcal{M}_1 can be directly characterized as those $\mu \in \mathcal{M}_0$ that additionally satisfy a collection of dual inequality constraints (see Appendix D.2). We will show that given the assumptions made on the transition matrix, a very small number of dual constraints suffice to pin down \mathcal{M}_1 . The binding constraints include:

1. 9 type-0 constraints: $\mu(s, a) \geq 0, \forall s, a$.
2. Constraint type-I-1:

$$\mu(1, 3)(p_{12} - p_{32}) \leq \mu(2, 1) + \mu(2, 3)(1 + p_{32} - p_{22}). \quad (33)$$

We use $\mu(1, 3)$ to represent $\mu(s_1, a_3)$, which should cause no confusion. This inequality is

³⁴When τ is positive, the agent can simply report as a function of s^{-1}, a^{-1} and s . Otherwise he has to condition on a^{-k}, \dots, a^{-1}, s .

³⁵We should point out that this does not imply $\mathcal{M}_1 \subset \overline{\mathcal{M}}_*$, because $\overline{\mathcal{M}}_*$ is not necessarily convex due to the non-linear equation (31). Actually we do not know whether convexity or the set inclusion holds.

derived from the following calculation:

$$\begin{aligned}
\mu(2, 2) &= \nu(1, 2, 2) + \nu(2, 2, 2) + \nu(3, 2, 2) \\
&\leq \sum_a \nu(1, 2, a) + \sum_a \nu(2, 2, a) + \sum_s \nu(3, s, 2) \\
&= \mu(1, 1)p_{12} + \mu(2, 1)p_{22} + \mu(3, 1)p_{32} + \mu(1, 2)p_{12} + \mu(2, 2)p_{22} + \mu(3, 2)p_{32} + \lambda_3 p_{32} \\
&= (\lambda_1 - \mu(1, 3))p_{12} + (\lambda_2 - \mu(2, 3))p_{22} + (\mu(1, 3) + \mu(2, 3) + \lambda_3)p_{32} \\
&= \lambda_2 + \mu(1, 3)(p_{32} - p_{12}) + \mu(2, 3)(p_{32} - p_{22}) \\
&= \mu(2, 2) + \mu(2, 1) + \mu(2, 3) + \mu(1, 3)(p_{32} - p_{12}) + \mu(2, 3)(p_{32} - p_{22}).
\end{aligned}$$

In the above derivation, we wrote $\nu(1, 2, 2)$ as a shorthand for $\nu(a_1, s_2, a_2)$.³⁶ The equalities follow from the fact that μ has the correct marginals λ .

This constraint binds if and only if $\nu(1, 2, 1) = \nu(1, 2, 3) = \nu(2, 2, 1) = \nu(2, 2, 3) = \nu(3, 1, 2) = \nu(3, 3, 2) = 0$. We present a figure to show the position of zeros in ν implied by this constraint:

$$\nu = \begin{pmatrix} + & + & + \\ 0 & + & 0 \\ + & + & + \end{pmatrix} \begin{pmatrix} + & + & + \\ 0 & + & 0 \\ + & + & + \end{pmatrix} \begin{pmatrix} + & 0 & + \\ + & + & + \\ + & 0 & + \end{pmatrix}$$

3. Constraint type-I-2:

$$\mu(1, 2)(p_{13} - p_{23}) \leq \mu(3, 1) + \mu(3, 2)(1 + p_{23} - p_{33}). \quad (34)$$

This is the same as the previous case, except that states s_2 and s_3 have switched (reports as well).

4. Constraint type-I-3:

$$\mu(2, 3)(p_{21} - p_{31}) \leq \mu(1, 2) + \mu(1, 3)(1 + p_{31} - p_{11}). \quad (35)$$

This can be obtained from case 2 by switching the states s_1 and s_2 (reports as well).

³⁶We also employed the identity $\mu(3, 1) + \mu(3, 2) = \mu(1, 3) + \mu(2, 3) = \lambda_3 - \mu(3, 3)$. This equation and its analogue for states 1 and 2 will be repeatedly used below.

We can use this subset of dual constraints to define a polytope \mathcal{Q} :

Definition 9 \mathcal{Q} is the set of $\mu \in \mathcal{M}_0$ that satisfy the three type-I constraints (33), (34), (35) listed above.

Since we are considering a smaller set of constraints, $\mathcal{Q} \supset \mathcal{M}_1$. We will however show that:

Lemma 17 \mathcal{Q} is a polytope with 22 vertices, each of which belongs to $\overline{\mathcal{M}_*}$.

If true, Lemma 17 together with Lemma 16 and $\mathcal{Q} \supset \mathcal{M}_1$ implies that $\mathcal{Q} = \mathcal{M}_1$. Proposition 3' will follow as a simple corollary. Hence, we just need to prove Lemma 17. By Lemma 2, \mathcal{Q} is a 4-dimensional polytope. Thus at any vertex μ of \mathcal{Q} , at least 4 dual constraints are binding. We can further show:

Lemma 18 If two type-I constraints above simultaneously bind at $\mu \in \mathcal{Q}$, then $\mu = \mu^{tt}$.

Proof. First suppose equations (33) and (34) hold simultaneously. The terms on the L.H.S. sum to at most $\frac{1}{2}(\mu(1,2) + \mu(1,3))$, while the terms on the R.H.S. sum to at least $\mu(2,1) + \mu(3,1) + \frac{1}{2}(\mu(2,3) + \mu(3,2)) = \mu(1,2) + \mu(1,3) + \frac{1}{2}(\mu(2,3) + \mu(3,2))$. Equality thus forces $\mu(1,2) = \mu(1,3) = \mu(2,3) = \mu(3,2) = 0$, implying $\mu = \mu^{tt}$. Essentially the same argument shows that (33) and (35) are not compatible unless $\mu = \mu^{tt}$. Finally for (34) and (35), we can use a similar argument and use $\mu(2,3) \leq \mu(3,1) + \mu(3,2)$. ■

By this lemma, at any vertex $\mu \neq \mu^{tt}$ of \mathcal{Q} , either μ satisfies one of the type-I constraints and is zero in three positions, or μ is zero in four positions. In the next few subsections, we carefully investigate each of these possibilities.

E.3 Constraint type-I binds (12 vertices)

Throughout this subsection (until the very end) we assume that constraint type-I-1 binds at $\mu \in \mathcal{Q}$, and $\mu \neq \mu^{tt}$. From equation (33), we have $\mu(1,3)(p_{12}-p_{32}) = \mu(2,1) + \mu(2,3)(1+p_{32}-p_{22})$.

Let us consider which 3 entries of μ could be zero. First we can rule out $\mu(1,3)$, because that would imply $\mu(2,1) = \mu(2,3) = 0$, and so $\mu = \mu^{tt}$. Next we rule out $\mu(3,1)$, because that would imply $\mu(1,3) \leq \mu(1,3) + \mu(2,3) = \mu(3,1) + \mu(3,2) = \mu(3,2) \leq \mu(1,2) + \mu(3,2) = \mu(2,1) + \mu(2,3)$. That contradicts the above equation. Similarly we can rule out $\mu(2,2)$, because that would also imply $\mu(2,1) + \mu(2,3) = \lambda_2 > \lambda_3 \geq \mu(1,3)$. It is rather obvious from the assumptions on p that $\lambda_1 > \lambda_2 > \lambda_3$. Finally we claim that $\mu(1,1)$ cannot be zero either. If $\mu(1,1) = 0$, then

$\mu(2, 1) = \mu(1, 1) + \mu(2, 1) + \mu(3, 1) - \mu(3, 1) \geq \lambda_1 - \lambda_3 > \lambda_3(p_{12} - p_{32}) \geq \mu(1, 3)(p_{12} - p_{32})$, which is a contradiction. The last inequality follows from

$$\begin{aligned} \lambda_1 - \lambda_3 &= \lambda_1(p_{11} - p_{13}) + \lambda_2(p_{21} - p_{23}) + \lambda_3(p_{31} - p_{33}) \\ &\geq \lambda_1(p_{11} - p_{13}) + \lambda_2(p_{21} - p_{23}) + \lambda_1(p_{31} - p_{33}) \\ &\geq \lambda_2(p_{21} - p_{23}) > \lambda_3(p_{12} - p_{32}). \end{aligned}$$

By this analysis, the 3 zeros in μ are chosen from $\mu(2, 1), \mu(2, 3), \mu(1, 2), \mu(3, 2)$ and $\mu(3, 3)$. However $\mu(2, 1), \mu(2, 3)$ cannot both be zero, because that would imply $\mu(1, 2) = \mu(3, 2) = 0$, $\mu(1, 3) = 0$ and $\mu = \mu^{tt}$. Similarly $\mu(1, 2), \mu(3, 2)$ cannot both be zero. Hence, we only have 4 possible choices, each involving one of $\mu(2, 1), \mu(2, 3)$ being zero, one of $\mu(1, 2), \mu(3, 2)$ being zero and $\mu(3, 3) = 0$. Together with equation (33), μ can be uniquely pinned down for each of these choices.

E.3.1 $\mu(2, 3) = \mu(3, 2) = \mu(3, 3) = 0$

We can obtain:

$$\mu(s, a) = \begin{pmatrix} \lambda_1 - (1 + p_{12} - p_{32})\lambda_3 & (p_{12} - p_{32})\lambda_3 & \lambda_3 \\ (p_{12} - p_{32})\lambda_3 & \lambda_2 - (p_{12} - p_{32})\lambda_3 & 0 \\ \lambda_3 & 0 & 0 \end{pmatrix}$$

The entries of μ are clearly all positive. Moreover, given the type-I-1 constraint and the type-zero constraints, any associated $\nu(a^{-1}, s, a)$ that solves the linear system defining \mathcal{M}_1 must have zeros in the following entries:

$$\nu = \begin{pmatrix} + & + & + \\ 0 & + & 0 \\ + & 0 & 0 \end{pmatrix} \begin{pmatrix} + & + & + \\ 0 & + & 0 \\ + & 0 & 0 \end{pmatrix} \begin{pmatrix} + & 0 & + \\ + & + & 0 \\ + & 0 & 0 \end{pmatrix}$$

To show $\mu \in \overline{\mathcal{M}_*}$, we need to demonstrate some τ and ν that satisfy (27) to (32). It will be convenient to work with the conditional probabilities $\tau(s, a|s^{-1}, a^{-1})$, for $\mu(s^{-1}, a^{-1}) > 0$. Since $\nu(a^{-1}, s, a) = \sum_{s^{-1}} \tau(s^{-1}, a^{-1}, s, a)$, we set $\tau(s, a|s^{-1}, a^{-1}) = 0$ whenever $\nu(a^{-1}, s, a) = 0$, which we can read from the preceding figure. Furthermore, the conditions (28) and (29) require that for fixed s^{-1}, a^{-1} , the $|S| \times |A|$ matrix $\tau(s, a|s^{-1}, a^{-1})$ has row sums $p(s|s^{-1})$ and column sums

$p(a|a^{-1})$. This uniquely determines $\tau(s, a|s^{-1}, a^{-1})$ as follows:

<table border="1" style="border-collapse: collapse; width: 100px; height: 100px;"> <tr><td>$p_{11} - p_{13}$</td><td>0</td><td>p_{13}</td></tr> <tr><td>0</td><td>p_{12}</td><td>0</td></tr> <tr><td>p_{13}</td><td>0</td><td>0</td></tr> </table>	$p_{11} - p_{13}$	0	p_{13}	0	p_{12}	0	p_{13}	0	0	<table border="1" style="border-collapse: collapse; width: 100px; height: 100px;"> <tr><td>$p_{11} - p_{23}$</td><td>$p_{12} - p_{22}$</td><td>p_{13}</td></tr> <tr><td>0</td><td>p_{22}</td><td>0</td></tr> <tr><td>p_{23}</td><td>0</td><td>0</td></tr> </table>	$p_{11} - p_{23}$	$p_{12} - p_{22}$	p_{13}	0	p_{22}	0	p_{23}	0	0	<table border="1" style="border-collapse: collapse; width: 100px; height: 100px;"> <tr><td>$p_{11} - p_{33}$</td><td>$p_{12} - p_{32}$</td><td>p_{13}</td></tr> <tr><td>0</td><td>p_{32}</td><td>0</td></tr> <tr><td>p_{33}</td><td>0</td><td>0</td></tr> </table>	$p_{11} - p_{33}$	$p_{12} - p_{32}$	p_{13}	0	p_{32}	0	p_{33}	0	0
$p_{11} - p_{13}$	0	p_{13}																											
0	p_{12}	0																											
p_{13}	0	0																											
$p_{11} - p_{23}$	$p_{12} - p_{22}$	p_{13}																											
0	p_{22}	0																											
p_{23}	0	0																											
$p_{11} - p_{33}$	$p_{12} - p_{32}$	p_{13}																											
0	p_{32}	0																											
p_{33}	0	0																											
$s^{-1}=1, a^{-1}=1$	$s^{-1}=2, a^{-1}=1$	$s^{-1}=3, a^{-1}=1$																											
<table border="1" style="border-collapse: collapse; width: 100px; height: 100px;"> <tr><td>$p_{21} - p_{13}$</td><td>$p_{22} - p_{12}$</td><td>p_{23}</td></tr> <tr><td>0</td><td>p_{12}</td><td>0</td></tr> <tr><td>p_{13}</td><td>0</td><td>0</td></tr> </table>	$p_{21} - p_{13}$	$p_{22} - p_{12}$	p_{23}	0	p_{12}	0	p_{13}	0	0	<table border="1" style="border-collapse: collapse; width: 100px; height: 100px;"> <tr><td>$p_{21} - p_{23}$</td><td>0</td><td>p_{23}</td></tr> <tr><td>0</td><td>p_{22}</td><td>0</td></tr> <tr><td>p_{23}</td><td>0</td><td>0</td></tr> </table>	$p_{21} - p_{23}$	0	p_{23}	0	p_{22}	0	p_{23}	0	0	<table border="1" style="border-collapse: collapse; width: 100px; height: 100px;"> <tr><td>$p_{11} - p_{33}$</td><td>0</td><td>p_{33}</td></tr> <tr><td>$p_{12} - p_{32}$</td><td>p_{32}</td><td>0</td></tr> <tr><td>p_{13}</td><td>0</td><td>0</td></tr> </table>	$p_{11} - p_{33}$	0	p_{33}	$p_{12} - p_{32}$	p_{32}	0	p_{13}	0	0
$p_{21} - p_{13}$	$p_{22} - p_{12}$	p_{23}																											
0	p_{12}	0																											
p_{13}	0	0																											
$p_{21} - p_{23}$	0	p_{23}																											
0	p_{22}	0																											
p_{23}	0	0																											
$p_{11} - p_{33}$	0	p_{33}																											
$p_{12} - p_{32}$	p_{32}	0																											
p_{13}	0	0																											
$s^{-1}=1, a^{-1}=2$	$s^{-1}=2, a^{-1}=2$	$s^{-1}=1, a^{-1}=3$																											

We note that condition (27) is obviously satisfied for $(s, a) = (2, 3), (3, 2), (3, 3)$. It also holds for $(s, a) = (2, 1)$, as $\mu(2, 1) = (p_{12} - p_{32})\mu(1, 3)$. Thus it holds for all (s, a) because given conditions (28) and (29), there are only 4 linearly-independent constraints like (27).

It remains to show that with this τ , the resulting ν as defined by (30) and (31) is always positive. For this we induct on k , with the base case $k = 0$ being immediate from $\mu \geq 0$. To complete the induction step, we notice that the only ‘‘conditional probability’’ $\tau(s, a|s^{-1}, a^{-1})$ that is negative is $\tau(s = 1, a = 2|s^{-1} = 2, a^{-1} = 1) = p_{12} - p_{22}$. By (31), we see that whenever $(a^{-1}, s, a) \neq (1, 1, 2)$, $\nu(a^{-k}, \dots, a^{-1}, s, a) \geq 0$. For $a^{-1} = 1, s = 1, a = 2$, we have by using (31) twice that (assuming $k \geq 2$):

$$\begin{aligned}
& \nu(a^{-k}, \dots, a^{-1} = 1, s = 1, a = 2) \\
&= \nu(a^{-k}, \dots, a^{-2}, s^{-1} = 2, a^{-1} = 1)(p_{12} - p_{22}) + \nu(a^{-k}, \dots, a^{-2}, s^{-1} = 3, a^{-1} = 1)(p_{12} - p_{32}) \\
&= \sum_{s^{-2}} \nu(a^{-k}, \dots, s^{-2}, a^{-2}) \cdot ((p_{12} - p_{22})\tau(2, 1|s^{-2}, a^{-2}) + (p_{12} - p_{32})\tau(3, 1|s^{-2}, a^{-2})).
\end{aligned} \tag{36}$$

We claim that $(p_{12} - p_{22})\tau(2, 1|s^{-2}, a^{-2}) + (p_{12} - p_{32})\tau(3, 1|s^{-2}, a^{-2}) \geq 0$ for all s^{-2}, a^{-2} . When $(s^{-2}, a^{-2}) \neq (1, 3)$, we see from the above figure for τ that the first summand is zero, while the second summand is positive. When $(s^{-2}, a^{-2}) = (1, 3)$, we have $\tau(2, 1|s^{-2}, a^{-2}) = p_{12} - p_{32} > 0$ and $\tau(3, 1|s^{-2}, a^{-2}) = p_{13}$. The claim then reduces to $p_{12} - p_{22} + p_{13} \geq 0$, which is equivalent to the assumption $p_{11} + p_{22} \leq 1$.

By this claim and equation (36), we deduce $\nu(a^{-k}, \dots, a^{-1} = 1, s = 1, a = 2) \geq 0$ from induction hypothesis. The result also holds when $k = 1$, as $\mu(2, 1)(p_{12} - p_{22}) + \mu(3, 1)(p_{12} - p_{32}) \geq$

0. Hence we have shown this μ belongs to $\overline{\mathcal{M}}_*$.

E.3.2 $\mu(2, 1) = \mu(3, 2) = \mu(3, 3) = 0$

Here we have:

$$\mu(s, a) = \begin{pmatrix} \lambda_1 - \lambda_3 & \frac{p_{12}-p_{32}}{1+p_{12}-p_{22}} \cdot \lambda_3 & \frac{1+p_{32}-p_{22}}{1+p_{12}-p_{22}} \cdot \lambda_3 \\ 0 & \lambda_2 - \frac{p_{12}-p_{32}}{1+p_{12}-p_{22}} \cdot \lambda_3 & \frac{p_{12}-p_{32}}{1+p_{12}-p_{22}} \cdot \lambda_3 \\ \lambda_3 & 0 & 0 \end{pmatrix}$$

Similar analysis yields a unique candidate $\tau(s, a|s^{-1}, a^{-1})$:

$p_{11} - p_{13}$	0	p_{13}
0	p_{12}	0
p_{13}	0	0

$s^{-1}=1, a^{-1}=1$

$p_{11} - p_{33}$	$p_{12} - p_{32}$	p_{13}
0	p_{32}	0
p_{33}	0	0

$s^{-1}=3, a^{-1}=1$

$p_{21} - p_{13}$	$p_{22} - p_{12}$	p_{23}
0	p_{12}	0
p_{13}	0	0

$s^{-1}=1, a^{-1}=2$

$p_{21} - p_{23}$	0	p_{23}
0	p_{22}	0
p_{23}	0	0

$s^{-1}=2, a^{-1}=2$

$p_{31} - p_{13}$	0	$p_{32} + p_{33} - p_{12}$
0	p_{32}	$p_{12} - p_{32}$
p_{13}	0	0

$s^{-1}=1, a^{-1}=3$

$p_{31} - p_{23}$	0	$p_{32} + p_{33} - p_{22}$
0	p_{32}	$p_{22} - p_{32}$
p_{23}	0	0

$s^{-1}=2, a^{-1}=3$

Since all the above entries in τ are positive, $\mu \in \overline{\mathcal{M}}_*$ as desired.

E.3.3 $\mu(2, 1) = \mu(1, 2) = \mu(3, 3) = 0$

$$\mu(s, a) = \begin{pmatrix} \lambda_1 - \frac{1+p_{32}-p_{22}}{1+p_{12}-p_{22}} \cdot \lambda_3 & 0 & \frac{1+p_{32}-p_{22}}{1+p_{12}-p_{22}} \cdot \lambda_3 \\ 0 & \lambda_2 - \frac{p_{12}-p_{32}}{1+p_{12}-p_{22}} \cdot \lambda_3 & \frac{p_{12}-p_{32}}{1+p_{12}-p_{22}} \cdot \lambda_3 \\ \frac{1+p_{32}-p_{22}}{1+p_{12}-p_{22}} \cdot \lambda_3 & \frac{p_{12}-p_{32}}{1+p_{12}-p_{22}} \cdot \lambda_3 & 0 \end{pmatrix}$$

Then $\tau(s, a|s^{-1}, a^{-1})$ can be filled out in the following way:

$p_{11} - p_{13}$	0	p_{13}
0	p_{12}	0
p_{13}	0	0

$s^{-1}=1, a^{-1}=1$

$p_{31} - p_{13}$	0	p_{13}
0	p_{32}	0
$p_{32} + p_{33} - p_{12}$	$p_{12} - p_{32}$	0

$s^{-1}=3, a^{-1}=1$

$p_{21} - p_{23}$	0	p_{23}
0	p_{22}	0
p_{23}	0	0

$s^{-1}=2, a^{-1}=2$

$p_{31} - p_{23}$	0	p_{23}
0	p_{32}	0
$p_{21} + p_{23} - p_{31}$	$p_{22} - p_{32}$	0

$s^{-1}=3, a^{-1}=2$

$p_{31} - p_{13}$	0	$p_{32} + p_{33} - p_{12}$
0	p_{32}	$p_{12} - p_{32}$
p_{13}	0	0

$s^{-1}=1, a^{-1}=3$

$p_{31} - p_{23}$	0	$p_{32} + p_{33} - p_{22}$
0	p_{32}	$p_{22} - p_{32}$
p_{23}	0	0

$s^{-1}=2, a^{-1}=3$

Again all the entries in τ are positive, so $\mu \in \overline{\mathcal{M}}_*$.

E.3.4 $\mu(2, 3) = \mu(1, 2) = \mu(3, 3) = 0$

In this case

$$\mu(s, a) = \begin{pmatrix} \lambda_1 - \lambda_3 & 0 & \lambda_3 \\ (p_{12} - p_{32})\lambda_3 & \lambda_2 - (p_{12} - p_{32})\lambda_3 & 0 \\ (1 - p_{12} + p_{32})\lambda_3 & (p_{12} - p_{32})\lambda_3 & 0 \end{pmatrix}$$

And τ is given by:

$p_{11} - p_{13}$	0	p_{13}
0	p_{12}	0
p_{13}	0	0

$s^{-1}=1, a^{-1}=1$

$p_{21} - p_{13}$	0	p_{13}
0	p_{22}	0
$p_{22} + p_{23} - p_{12}$	$p_{12} - p_{22}$	0

$s^{-1}=2, a^{-1}=1$

$p_{31} - p_{13}$	0	p_{13}
0	p_{32}	0
$p_{32} + p_{33} - p_{12}$	$p_{12} - p_{32}$	0

$s^{-1}=3, a^{-1}=1$

$p_{21} - p_{23}$	0	p_{23}
0	p_{22}	0
p_{23}	0	0

$s^{-1}=2, a^{-1}=2$

$p_{31} - p_{23}$	0	p_{23}
0	p_{32}	0
$p_{21} + p_{23} - p_{31}$	$p_{22} - p_{32}$	0

$s^{-1}=3, a^{-1}=2$

$p_{11} - p_{33}$	0	p_{33}
$p_{12} - p_{32}$	p_{32}	0
p_{13}	0	0

$s^{-1}=1, a^{-1}=3$

The only negative entry in τ is that $\tau(s = 3, a = 2 | s^{-1} = 2, a^{-1} = 1) = p_{12} - p_{22}$. But we can still inductively prove the positivity of ν by establishing an analogue of equation (36). The key is that $(p_{12} - p_{22})\tau(2, 1 | s^{-2}, a^{-2}) + (p_{12} - p_{32})\tau(3, 1 | s^{-2}, a^{-2}) \geq 0$ for all s^{-2}, a^{-2} . Hence this μ belongs to $\overline{\mathcal{M}}_*$ as well.

This completes the proof of the crucial Lemma 17 when the vertex μ is such that the dual constraint type-I-1 binds. Due to the fact that the other two type-I constraints can be obtained by permuting states and reports, an almost identical argument applies to those cases as well. We have thus found 12 vertices of \mathcal{Q} , which all belong to $\overline{\mathcal{M}}_*$.

E.4 μ has 4 zeros

Given the assumptions on the transition matrix, we have $\frac{1}{2} > \lambda_1 > \lambda_2 > \lambda_3 > 0$. Thus, there are exactly 10 such $\mu \in \mathcal{M}_0$ that are zero in four positions:

$$\begin{array}{cc}
 \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix} & \begin{pmatrix} 0 & \lambda_1 - \lambda_3 & \lambda_3 \\ \lambda_2 & 0 & 0 \\ \lambda_1 - \lambda_2 & \lambda_2 + \lambda_3 - \lambda_1 & 0 \end{pmatrix} \\
 \begin{pmatrix} 0 & \lambda_2 & \lambda_1 - \lambda_2 \\ \lambda_2 & 0 & 0 \\ \lambda_1 - \lambda_2 & 0 & \lambda_2 + \lambda_3 - \lambda_1 \end{pmatrix} & \begin{pmatrix} 0 & \lambda_2 & \lambda_1 - \lambda_2 \\ \lambda_1 - \lambda_3 & 0 & \lambda_2 + \lambda_3 - \lambda_1 \\ \lambda_3 & 0 & 0 \end{pmatrix} \\
 \begin{pmatrix} 0 & \lambda_1 - \lambda_3 & \lambda_3 \\ \lambda_1 - \lambda_3 & \lambda_2 + \lambda_3 - \lambda_1 & 0 \\ \lambda_3 & 0 & 0 \end{pmatrix} & \begin{pmatrix} \lambda_1 - \lambda_3 & 0 & \lambda_3 \\ \lambda_3 & \lambda_2 - \lambda_3 & 0 \\ 0 & \lambda_3 & 0 \end{pmatrix} \\
 \begin{pmatrix} \lambda_1 - \lambda_3 & \lambda_3 & 0 \\ 0 & \lambda_2 - \lambda_3 & \lambda_3 \\ \lambda_3 & 0 & 0 \end{pmatrix} & \begin{pmatrix} \lambda_1 - \lambda_2 & \lambda_2 & 0 \\ \lambda_2 - \lambda_3 & 0 & \lambda_3 \\ \lambda_3 & 0 & 0 \end{pmatrix} \\
 \begin{pmatrix} \lambda_1 + \lambda_3 - \lambda_2 & \lambda_2 - \lambda_3 & 0 \\ \lambda_2 - \lambda_3 & 0 & \lambda_3 \\ 0 & \lambda_3 & 0 \end{pmatrix} & \begin{pmatrix} \lambda_1 + \lambda_3 - \lambda_2 & \lambda_2 - \lambda_3 & 0 \\ \lambda_2 - \lambda_3 & 0 & \lambda_3 \\ 0 & \lambda_3 & 0 \end{pmatrix}
 \end{array}$$

For each μ listed above, τ can be uniquely determined as before. Without laying out further details, let us simply note that all these τ are everywhere positive, so $\mu \in \overline{\mathcal{M}}_*$ as expected. This completes the proof of Lemma 17 and thus of Proposition ?? . ■

F Insufficiency of \mathcal{M}_k

F.1 Insufficiency of \mathcal{M}_1

Proposition 4 *Let $|S| = 3$. Fix a neighborhood of the transition matrix for (s_n) i.i.d. and uniform. Then this neighborhood contains two open sets of transition matrices, one with $\mathcal{M}_1 = \mathcal{M}_k$ for all k , and one with $\mathcal{M}_1 \neq \mathcal{M}_k$ for some k .*

The first half of the statement follows directly from Proposition ???. Here we prove the second half via the following lemma:

Lemma 19 *Assume $p_{11} > p_{33} > p_{22} > \max\{p_{21}, p_{31}\} > \max\{p_{12}, p_{32}\} > \max\{p_{13}, p_{23}\} > 0$, $p_{11} + p_{22} < 1$ and further $p_{33} - p_{22} < (p_{22} - p_{23})^2$. Then there exists k such that $\mathcal{M}_1 \neq \mathcal{M}_k$.*

Proof. We should first point out that the assumptions are satisfied by:

$$p_\varepsilon = \begin{pmatrix} \frac{1}{3} + 5\varepsilon & \frac{1}{3} - 2\varepsilon & \frac{1}{3} - 3\varepsilon \\ \frac{1}{3} + \varepsilon & \frac{1}{3} + 3\varepsilon & \frac{1}{3} - 4\varepsilon \\ \frac{1}{3} - \varepsilon & \frac{1}{3} - 2\varepsilon - \varepsilon^2 & \frac{1}{3} + 3\varepsilon + \varepsilon^2 \end{pmatrix}$$

with ε small. Since these assumptions define an open set of transition matrices, the current lemma does imply Proposition 4. It is also worth noting that the assumptions here differ from those in Proposition ??? only in that the comparison between p_{22} and p_{33} is reversed. This again highlights the observation that the equality $\mathcal{M}_1 = \mathcal{M}_k$ depends on very fine details of the process.

Now we present the proof of the lemma, starting with a technical observation:

Claim 4 $\lambda_1 > \lambda_3$ and $\lambda_2 - \lambda_3 > \frac{p_{33} - p_{22}}{p_{22} - p_{23}} \cdot \lambda_3$.

Proof. We note that up to proportionality, $\lambda_1 = p_{21}p_{31} + p_{23}p_{31} + p_{32}p_{21}$, $\lambda_2 = p_{12}p_{32} + p_{13}p_{32} + p_{31}p_{12}$ and $\lambda_3 = p_{13}p_{23} + p_{12}p_{23} + p_{21}p_{13}$. Thus we immediately have $\lambda_1 > \lambda_3$. We also have

$$\begin{aligned} \lambda_2 &= p_{12}p_{33} + p_{12}(p_{32} - p_{33}) + p_{12}p_{32} - (p_{12} - p_{13})p_{32} + p_{12}p_{31} \\ &= p_{12} + p_{12}(p_{32} - p_{33}) - (p_{12} - p_{13})p_{32}. \end{aligned}$$

$$\begin{aligned} \lambda_3 &= p_{12}p_{23} - (p_{12} - p_{13})p_{23} + p_{12}p_{23} + p_{12}p_{21} - (p_{12} - p_{13})p_{21} \\ &= p_{12} - p_{12}(p_{22} - p_{23}) - (p_{12} - p_{13}) \cdot (p_{21} + p_{23}). \end{aligned}$$

Thus $\lambda_2 - \lambda_3 > p_{12}(p_{22} - p_{23}) > \lambda_3(p_{22} - p_{23}) > \frac{p_{33} - p_{22}}{p_{22} - p_{23}} \cdot \lambda_3$ by assumption. ■

To prove the lemma, we show that the following μ belongs to \mathcal{M}_1 but not to every \mathcal{M}_k :

$$\mu(s, a) = \begin{pmatrix} \lambda_1 - \lambda_3 & 0 & \lambda_3 \\ \lambda_3 & \lambda_2 - \lambda_3 & 0 \\ 0 & \lambda_3 & 0 \end{pmatrix}$$

Consider the following $\nu(a^{-1}, s, a)$:

$$\begin{aligned} \nu(a^{-1} = 1, \cdot, \cdot) &= \begin{pmatrix} (\lambda_1 - \lambda_3)(p_{11} - p_{13}) + \lambda_3(p_{21} - p_{13}) & 0 & \lambda_1 p_{13} \\ (\lambda_1 - \lambda_3)p_{13} + \lambda_3(p_{11} + p_{13} - p_{21}) & (\lambda_1 - \lambda_3)(p_{12} - p_{13}) + \lambda_3(p_{12} - p_{23}) & 0 \\ 0 & (\lambda_1 - \lambda_3)p_{13} + \lambda_3 p_{23} & 0 \end{pmatrix} \\ \nu(a^{-1} = 2, \cdot, \cdot) &= \begin{pmatrix} (\lambda_2 - \lambda_3)(p_{21} - p_{23}) + \lambda_3(p_{31} - p_{23}) & 0 & \lambda_2 p_{23} \\ (\lambda_2 - \lambda_3)p_{23} + \lambda_3(p_{21} + p_{23} - p_{31}) & (\lambda_2 - \lambda_3)(p_{22} - p_{23}) - \lambda_3(p_{33} - p_{22}) & 0 \\ 0 & (\lambda_2 - \lambda_3)p_{23} + \lambda_3 p_{33} & 0 \end{pmatrix} \\ \nu(a^{-1} = 3, \cdot, \cdot) &= \begin{pmatrix} \lambda_3(p_{11} - p_{33}) & 0 & \lambda_3 p_{33} \\ \lambda_3(p_{31} + p_{33} - p_{11}) & \lambda_3(p_{32} - p_{13}) & 0 \\ 0 & \lambda_3 p_{13} & 0 \end{pmatrix} \end{aligned}$$

Note that ν is positive: the only issue is $\nu(2, 2, 2)$, which is positive by the preceding claim. In fact, we can show it is the only positive $\nu(a^{-1}, s, a)$ solving the linear system (2) to (4). This is because by equation (2), $\nu(a^{-1}, s, a) = 0$ whenever $\mu(s, a) = 0$. The remaining entries of ν can then be uniquely solved by (3) and (4).

We can in fact apply this argument to the iterative linear system (17) to (19), to obtain a unique candidate $\nu(a^{-k}, \dots, a^{-1}, s, a)$ that solves the linear system defining \mathcal{M}_k (assuming $\mu \in \mathcal{M}_k$). Let us focus on those ν with $a^{-k} = \dots = a^{-1} = 2$. We have $\nu(\dots, s, a) = 0$ when $\mu(s, a) = 0$. Then by equation (19) we obtain:

$$\begin{aligned} \nu(a^{-k} = \dots = a^{-1} = 2, s = 3, a = 2) &= \sum_a \nu(a^{-k} = \dots = a^{-1} = 2, s = 3, a) \\ &= \nu(a^{-k} = \dots = a^{-2} = 2, s^{-1} = 2, a^{-1} = 2) \cdot p_{23} + \nu(a^{-k} = \dots = a^{-2} = 2, s^{-1} = 3, a^{-1} = 2) \cdot p_{33}. \end{aligned} \tag{37}$$

Using this and $\nu(\dots, s = 1, a = 2) = 0$, we can then solve for $\nu(\dots, s = 2, a = 2)$ from equation (18):

$$\begin{aligned} &\nu(a^{-k} = \dots = a^{-1} = 2, s = 2, a = 2) \\ &= \sum_s \nu(a^{-k} = \dots = a^{-1} = 2, s, a = 2) - \nu(a^{-k} = \dots = a^{-1} = 2, s = 3, a = 2) \\ &= \nu(a^{-k} = \dots = 2, s^{-1} = 2, a^{-1} = 2)(p_{22} - p_{23}) + \nu(a^{-k} = \dots = 2, s^{-1} = 3, a^{-1} = 2)(p_{22} - p_{33}). \end{aligned} \tag{38}$$

In the above calculation we have used the following identity:

$$\begin{aligned}
& \sum_s \nu(a^{-k} = \dots = a^{-1} = 2, s, a = 2) = \lambda(a^{-k} = \dots = a = 2) \\
& = \lambda(a^{-k} = \dots = a^{-1} = 2) \cdot p_{22} \\
& = \sum_{s^{-1}} \nu(a^{-k} = \dots = a^{-2} = 2, s^{-1}, a^{-1} = 2) \cdot p_{22}.
\end{aligned}$$

If we now define $x_k = \nu(a^{-k} = \dots = 2, s = 2, a = 2)$ and $y_k = \nu(a^{-k} = \dots = 2, s = 3, a = 2)$, then equations (37) and (38) reduce to the following linear recurrence relation:³⁷

$$\begin{bmatrix} x_k \\ y_k \end{bmatrix} = \begin{bmatrix} p_{22} - p_{23} & p_{22} - p_{33} \\ p_{23} & p_{33} \end{bmatrix} \cdot \begin{bmatrix} x_{k-1} \\ y_{k-1} \end{bmatrix} \quad (39)$$

The matrix above has two positive eigenvalues p_{22} and $p_{33} - p_{23}$. The associated eigenvectors are $\begin{bmatrix} \frac{p_{33}-p_{22}}{p_{23}} \\ -1 \end{bmatrix}$ and $\begin{bmatrix} 1 \\ -1 \end{bmatrix}$. The initial vector $\begin{bmatrix} x_0 \\ y_0 \end{bmatrix}$ has positive coordinates, which is not an

eigenvector. By standard results in linear algebra, if $p_{22} > p_{33} - p_{23}$ then $\begin{bmatrix} x_k \\ y_k \end{bmatrix} \rightarrow \begin{bmatrix} \frac{p_{33}-p_{22}}{p_{23}} \\ -1 \end{bmatrix}$ up

to proportionality. If instead $p_{22} < p_{33} - p_{23}$, then $\begin{bmatrix} x_k \\ y_k \end{bmatrix} \rightarrow \begin{bmatrix} 1 \\ -1 \end{bmatrix}$ up to proportionality. In either case, we cannot always have $x_k, y_k \geq 0$. This implies $\mu \notin \mathcal{M}_k$ for some k .

In the knife-edge case $p_{22} = p_{33} - p_{23}$, we obtain from (39) and induction that:

$$\begin{bmatrix} x_k \\ y_k \end{bmatrix} = \begin{bmatrix} p_{22}^k - k \cdot p_{23} \cdot p_{22}^{k-1} & -k \cdot p_{23} \cdot p_{22}^{k-1} \\ k \cdot p_{23} \cdot p_{22}^{k-1} & p_{22}^k + k \cdot p_{23} \cdot p_{22}^{k-1} \end{bmatrix} \cdot \begin{bmatrix} x_0 \\ y_0 \end{bmatrix}$$

Thus for sufficiently large k , $x_k = (p_{22}^k - k \cdot p_{23} \cdot p_{22}^{k-1}) \cdot x_0 + (-k \cdot p_{23} \cdot p_{22}^{k-1}) \cdot y_0 < 0$, again implying that $\mu \notin \mathcal{M}_k$. ■

³⁷We note that when $\nu(\dots)$ can be uniquely solved from the iterative linear system, these recursive relations are exactly the same as equation (31), for some τ that is also uniquely determined. This is our original motivation for considering the set $\underline{\mathcal{M}}_*$.

F.2 Additional examples showing the insufficiency of \mathcal{M}_1

F.2.1 Pairs vs. triples

Our leading example in the main text shows that checking triples can be strictly better than checking pairs. However, we obtain this conclusion assuming that the designer only checks *consecutive pairs*. Suppose the designer could keep track of transition counts of both (a^{-1}, a) and (a^{-2}, a) . Is this as good as keeping track of the frequencies of triples (a^{-2}, a^{-1}, a) ? To formalize the question, we introduce a new set \mathcal{M}'_2 to capture the set of joint distributions that are undetectable given a designer who has memory 2 and checks pairs only:

Definition 10 \mathcal{M}'_2 is the set of distributions $\mu \in \Delta(S \times A)$ such that the following linear system in $\nu(a^{-2}, \dots, a^{-1}, s, a)$ has a solution:

$$\begin{aligned} \sum_{a^{-2}, a^{-1}} \nu(a^{-2}, a^{-1}, s, a) &= \mu(s, a). \\ \sum_{a^{-2}, s} \nu(a^{-2}, a^{-1}, s, a) &= \lambda(a^{-1})p(a|a^{-1}). \\ \sum_{a^{-1}, s} \nu(a^{-2}, a^{-1}, s, a) &= \lambda(a^{-2}) \cdot \left(\sum_{a^{-1}} p(a^{-1}|a^{-2})p(a|a^{-1}) \right). \\ \sum_a \nu(a^{-2}, a^{-1}, s, a) &= \sum_{s^{-1}} \left(p(s|s^{-1}) \cdot \sum_{a^{-3}} \nu(a^{-3}, a^{-2}, s^{-1}, a^{-1}) \right). \\ \nu(a^{-2}, a^{-1}, s, a) &\geq 0. \end{aligned}$$

Proposition 12 $\mathcal{M}'_2 \supset \mathcal{M}_2$ and there exists an open set of three-state processes for which $\mathcal{M}'_2 \neq \mathcal{M}_2$.

Proof. $\mathcal{M}'_2 \supset \mathcal{M}_2$ because the condition (3) defining \mathcal{M}_2 is stronger than the second and third equation above combined. For the second half, we consider the following “cyclic” process:³⁸

$$p = \begin{pmatrix} 0.65 & 0.3 & 0.05 \\ 0.05 & 0.65 & 0.3 \\ 0.3 & 0.05 & 0.65 \end{pmatrix}$$

³⁸For such cyclic processes, we are able to completely characterize the 57 vertices of \mathcal{M}_1 using dual constraints.

By direct numerical computation, the following μ belongs to \mathcal{M}'_2 but not \mathcal{M}_2 :

$$\mu(s, a) = \begin{pmatrix} \frac{39}{141} & 0 & \frac{8}{141} \\ \frac{8}{141} & \frac{15}{141} & \frac{24}{141} \\ 0 & \frac{32}{141} & \frac{15}{141} \end{pmatrix}$$

The result for an open set of processes follows from continuity (Lemma 2 and its analogue for \mathcal{M}'_2). ■

F.2.2 One more round

From the proof of Theorem 1, we know that $\mathcal{M}_0 = \mathcal{M}_k$ for every k if and only if $\mathcal{M}_0 = \mathcal{M}_1$. A natural question is whether in general, identifying some longer string does not help enables one to conclude that no longer memory helps. We answer this question in the negative by considering the following process:

$$p = \begin{pmatrix} 0.51 & 0.48 & 0.01 \\ 0.01 & 0.51 & 0.48 \\ 0.48 & 0.01 & 0.51 \end{pmatrix}$$

Numerical computation shows that for this process, $\mathcal{M}_1 = \mathcal{M}_2 \neq \mathcal{M}_3$. We do not know whether such examples are non-generic.

F.3 Insufficiency of \mathcal{M}_k

Here we show that the analysis for the leading example in Section 2 can be generalized to any memory k . Specifically, consider any Markov chain with transition matrix:

$$p = \begin{pmatrix} 1 - c_1 & c_1 & 0 \\ 0 & 1 - c_2 & c_2 \\ c_3 & 0 & 1 - c_3 \end{pmatrix}$$

with $c_3 \geq c_1 > c_2 > 0$. We establish the following result:

Claim 5 *For any such process p and any k , there exists $\ell > k$ such that $\mathcal{M}_k \neq \mathcal{M}_\ell$.*

Proof. Like in the main text, let us consider the joint distribution μ_x given by:

$$\mu(s, a) = \begin{pmatrix} \lambda_1 - x & x & 0 \\ 0 & \lambda_2 - x & x \\ x & 0 & \lambda_3 - x \end{pmatrix}$$

The same argument as in Footnote 14 shows that if $\mu_{x_k} \in \mathcal{M}_k$, then $x_k \rightarrow 0$. The claim will follow once we construct, for each k , some $x_k > 0$ such that $\mu_{x_k} \in \mathcal{M}_k$. Fix k and $x > 0$. Due to the many zeros in p and in μ_x , we can iteratively solve for the (unique) $\nu(a^{-k}, \dots, a^{-1}, s, a)$ that could make $\mu_x \in \mathcal{M}_k$. The result is the following:

$$\nu(a^{-k}, \dots, a^{-1}, s, a) = \begin{cases} x \cdot q(a^{-(k-1)}|a^{-k}) \cdots q(a|a^{-1}), & \text{if } s = a - 1; \\ \lambda(a^{-k})p(a^{-(k-1)}|a^{-k}) \cdots p(a|a^{-1}) - \nu(a^{-k}, \dots, a^{-1}, s - 1, a), & \text{if } s = a; \\ 0, & \text{if } s = a + 1 \end{cases}$$

where q represents the following transition matrix (not necessarily stochastic):

$$q = \begin{pmatrix} 1 - c_3 & c_1 & 0 \\ 0 & 1 - c_1 & c_2 \\ c_3 & 0 & 1 - c_2 \end{pmatrix}$$

It is straightforward to check that $\nu \geq 0$ for x sufficiently small. ■

Assuming that Proposition ?? holds, we can obtain the following stronger result:

Proposition 6 *Let $|S| = 3$ and fix $k \geq 1$. There exists an open set of transition matrices such that, for all matrices within that set, $\mathcal{M}_k \neq \mathcal{M}_\infty$, yet $\mathcal{C}_k = \mathcal{C}_\infty$.*

Proof. Let us pick c_1, c_2, c_3 sufficiently close to $\frac{1}{2}$, and fix the process p as in the previous claim. For any full-support process p_n sufficiently close to p , it holds by Proposition ?? that $\mathcal{C}_1(p_n) = \mathcal{C}_k(p_n) = \mathcal{C}_\infty(p_n)$. Furthermore, we know that $\mathcal{M}_k(p) \neq \mathcal{M}_l(p)$ for some $l > k$. By Lemma 2, we can conclude $\mathcal{M}_k(p_n) \neq \mathcal{M}_l(p_n)$ which implies the desired result. ■

G Sufficiency of \mathcal{C}_1

We have seen that whether longer memory restricts the set of undetectable joint distributions depends on fine details of the process. Despite this subtlety, we are going to show in this appendix that whether memory 1 suffices for implementation admits a clear-cut answer:

Proposition 5 *Let $|S| = 3$. For every $k \geq 1$, $\mathcal{C}_1 = \mathcal{C}_k$ if $0 < p(s' | s) \leq \beta = \frac{\sqrt{5}-1}{2}$ for all s, s' .*

G.1 Proof outline and the cone $\overline{\mathcal{C}}_*$

The argument we present follows a similar logic to the proof of $\mathcal{M}_1 = \mathcal{M}_k$, with some modifications. First, we will define a cone $\overline{\mathcal{C}}_*$ that plays the role of $\overline{\mathcal{M}}_*$, with the property that $\overline{\mathcal{C}}_* \subset \mathcal{C}_k$ for every k . We then show that when transition probabilities are not too big, every *extremal ray* in \mathcal{C}_1 belongs to $\overline{\mathcal{C}}_*$. While we were previously able to enumerate the vertices of \mathcal{M}_1 under very specific assumptions on p , this is no longer possible in the current setting. Instead we study general properties of any μ that belongs to an extremal ray. We carefully classify such μ into a dozen cases, and in each case we directly construct some τ that ensures $\mu \in \overline{\mathcal{C}}_*$. It is this casework that results in the length of the proof.

Let us begin by defining the cone $\overline{\mathcal{C}}_*$:

Definition 11 $\overline{\mathcal{C}}_*$ is the set of joint distributions $\mu \in \Delta(S \times A)$ for which the following system of equations admits a solution in τ and ν :

$$\tau(s^{-1}, a^{-1}, s, a) = 0, \quad \forall s^{-1} = a^{-1}, s \neq a. \quad (40)$$

$$\sum_{s^{-1}, a^{-1}} \tau(s^{-1}, a^{-1}, s, a) = \mu(s, a). \quad (41)$$

$$\sum_a \tau(s^{-1}, a^{-1}, s, a) = \mu(s^{-1}, a^{-1}) \cdot p(s|s^{-1}). \quad (42)$$

$$\sum_s \tau(s^{-1}, a^{-1}, s, a) = \mu(s^{-1}, a^{-1}) \cdot p(a|a^{-1}). \quad (43)$$

$$\nu(s, a) = \mu(s, a). \quad (44)$$

$$\nu(a^{-k}, \dots, a^{-1}, s, a) = \sum_{s^{-1}} \nu(a^{-k}, \dots, a^{-2}, s^{-1}, a^{-1}) \cdot \tau(s, a|s^{-1}, a^{-1}), \quad \forall k \geq 1. \quad (45)$$

$$\nu(a^{-k}, \dots, a^{-1}, s, a) \geq 0, \quad \forall s \neq a. \quad (46)$$

Conditions (41) to (45) are exactly the same as (27) to (31) in Definition 8. The agent reports according to a joint Markov chain on $S \times A$ such that both (s_n) and (a_n) could be viewed as autonomous. However, condition (46) is weaker than (32) in that we only require the resulting distribution ν to be positive *off the diagonal*. We make this relaxation because we will let the agent randomize between truth-telling and the reporting policy given by τ . Since truth-telling induces a distribution ν^{tt} that is strictly positive on the diagonal, some mixture between ν^{tt} and ν will be positive everywhere. The technical condition (40) ensures that such a mixture does result from a randomized policy, as we show in the following lemma:

Lemma 20 *When p has full support, $\overline{\mathcal{C}}_*$ is a cone centered at μ^{tt} and $\overline{\mathcal{C}}_* \subset \mathcal{C}_k$ for every k .*

Proof. Take $\mu \in \overline{\mathcal{C}}_*$. For any $\varepsilon > 0$ consider $\tilde{\mu} = \varepsilon\mu + (1 - \varepsilon)\mu^{tt}$. We claim that $\tilde{\mu} \in \overline{\mathcal{C}}_*$. For this we define $\tilde{\tau} = \varepsilon\tau + (1 - \varepsilon)\tau^{tt}$, where $\tau^{tt}(s^{-1}, a^{-1}, s, a) = \lambda(s^{-1})p(s|s^{-1})$ if $(a^{-1}, a) = (s^{-1}, s)$ and zero otherwise. Then $\tilde{\tau}$ and $\tilde{\mu}$ satisfy the linear conditions (40) to (43). We can similarly define $\tilde{\nu}^{tt}$ and $\tilde{\nu} = \varepsilon\nu + (1 - \varepsilon)\nu^{tt}$, so that (44) and (46) hold. It remains to check (45), which is typically non-linear. This is where condition (40) comes in. For $s \neq a$, we have $\tilde{\nu}(a^{-k}, \dots, a^{-1}, s, a) = \varepsilon\nu(a^{-k}, \dots, a^{-1}, s, a)$. Also $\tilde{\tau}(s, a|s^{-1}, a^{-1}) = \tau(s, a|s^{-1}, a^{-1})$ because they are both zero when $s^{-1} = a^{-1}$ and both given by $\frac{\tau(s^{-1}, a^{-1}, s, a)}{\mu(s^{-1}, a^{-1})}$ when $s^{-1} \neq a^{-1}$. Thus (45) holds when $s \neq a$. Using (42), we see that the sum of equation (45) over a is linear in ν . Thus for $\tilde{\mu}, \tilde{\tau}, \tilde{\nu}$, (45) holds for $s = a$ as well. This completes the proof that $\overline{\mathcal{C}}_*$ is a cone.

To prove $\overline{\mathcal{C}}_* \subset \mathcal{C}_k$, we note that if μ, τ satisfy equations (41) to (43), then μ and ν as defined by (44) and (45) is a (possibly negative) solution to the linear system defining \mathcal{M}_k . Take $\varepsilon > 0$, then $\tilde{\mu}$ and $\tilde{\nu}$ is also a solution. Condition (46) ensures that $\tilde{\nu}$ is positive whenever $s \neq a$. Since p has full-support, ν^{tt} is strictly positive when $s = a$, which implies that $\tilde{\nu}^{tt} = (1 - \varepsilon)\nu^{tt} + \varepsilon\nu$ is everywhere positive for sufficiently small ε . Hence $\tilde{\mu} = (1 - \varepsilon)\mu^{tt} + \varepsilon\mu \in \mathcal{M}_k$, and $\mu \in \mathcal{C}_k$ as desired. ■

With this lemma, we restate Proposition ?? in the following stronger form:

Proposition 5' *Suppose $|S| = 3$ and $0 < p(s'|s) \leq \beta = \frac{\sqrt{5}-1}{2}$. Take any vertex μ in \mathcal{M}_1 that is adjacent to μ^{tt} , then $\mu \in \overline{\mathcal{C}}_*$.*

We will alternatively speak of μ as lying on an “extremal ray” in \mathcal{C}_1 . Our analysis below does not depend on the specific choice of μ on such a ray, as we will only be concerned with off-diagonal values of μ, τ, ν and especially the ratios among them.

G.2 Binding dual constraints for adjacent μ

From Appendix D.2, we know that \mathcal{M}_1 consists of those $\mu \in \mathcal{M}_0$ that satisfy a finite collection of dual inequality constraints. Every vertex μ of \mathcal{M}_1 is determined by 5 marginal equalities as well as 4 binding dual inequalities. We can say more if μ is adjacent to μ^{tt} :

Lemma 21 *If $\mu \in \mathcal{M}_1$ is a vertex adjacent to μ^{tt} , then there are 3 linearly-independent dual inequalities that simultaneously bind at μ and μ^{tt} .*

Proof. If there are at most 2 dual inequalities binding at both μ and μ^{tt} , consider those $\Delta \in R^{|S| \times |A|}$ that have marginals zero and satisfy those dual constraints with equality. This is at least a 2-dimensional subspace, so we can take Δ to not be proportional to $\mu - \mu^{tt}$. For ε sufficiently small, consider $\mu' = .5\mu + .5\mu^{tt} + \varepsilon\Delta$ and $\mu'' = .5\mu + .5\mu^{tt} - \varepsilon\Delta$. Due to the construction of Δ , both μ' and μ'' satisfy the marginal equalities as well as dual inequalities. Thus μ' and μ'' both belong to \mathcal{M}_1 . But $\mu, \mu', \mu^{tt}, \mu''$ form a non-degenerate parallelogram, contradicting μ being a vertex adjacent to μ^{tt} . ■

To make use of this lemma, we now turn to analyzing which combinations of 3 dual constraints can occur.

G.2.1 Type-0 constraints

Claim 6 *The type-0 constraints that bind at μ^{tt} are $\mu(s, a) = 0$ for $s \neq a$.*

Proof. This is obvious. ■

G.2.2 Type-I constraints

Recall from Appendix D.2 that a type-I constraint corresponds to six entries in $\nu(a^{-1}, s, a)$ to be zero, see for example (24). In order for such a dual constraint to bind at μ^{tt} , these zero entries must be off-diagonal. It is not hard to see that there are 18 such constraints, 6 of which are shown below and others are obtained by permuting states and reports:

$$\nu = \begin{pmatrix} + & + & + \\ + & + & + \\ 0 & 0 & + \end{pmatrix} \begin{pmatrix} + & + & + \\ + & + & + \\ 0 & 0 & + \end{pmatrix} \begin{pmatrix} + & + & 0 \\ + & + & 0 \\ + & + & + \end{pmatrix} \quad (\text{A1})$$

$$\nu = \begin{pmatrix} + & + & 0 \\ + & + & 0 \\ + & + & + \end{pmatrix} \begin{pmatrix} + & + & 0 \\ + & + & 0 \\ + & + & + \end{pmatrix} \begin{pmatrix} + & + & + \\ + & + & + \\ 0 & 0 & + \end{pmatrix} \quad (\text{A2})$$

$$\nu = \begin{pmatrix} + & + & + \\ + & + & + \\ 0 & 0 & + \end{pmatrix} \begin{pmatrix} + & + & 0 \\ + & + & 0 \\ + & + & + \end{pmatrix} \begin{pmatrix} + & + & + \\ + & + & + \\ 0 & 0 & + \end{pmatrix} \quad (\text{A3})$$

$$\nu = \begin{pmatrix} + & + & + \\ + & + & + \\ 0 & 0 & + \end{pmatrix} \begin{pmatrix} + & + & 0 \\ + & + & 0 \\ + & + & + \end{pmatrix} \begin{pmatrix} + & + & 0 \\ + & + & 0 \\ + & + & + \end{pmatrix} \quad (\text{A4})$$

$$\nu = \begin{pmatrix} + & + & 0 \\ + & + & 0 \\ + & + & + \end{pmatrix} \begin{pmatrix} + & + & + \\ + & + & + \\ 0 & 0 & + \end{pmatrix} \begin{pmatrix} + & + & + \\ + & + & + \\ 0 & 0 & + \end{pmatrix} \quad (\text{A5})$$

$$\nu = \begin{pmatrix} + & + & 0 \\ + & + & 0 \\ + & + & + \end{pmatrix} \begin{pmatrix} + & + & + \\ + & + & + \\ 0 & 0 & + \end{pmatrix} \begin{pmatrix} + & + & 0 \\ + & + & 0 \\ + & + & + \end{pmatrix} \quad (\text{A6})$$

Claim 7 *There are 18 type-I constraints that bind at μ^{tt} , corresponding to the figures (A1)–(A6) and their permutations. Moreover, (A1) and (A2) cannot bind at μ unless $\mu(1, 3) = \mu(2, 3) = \mu(3, 1) = \mu(3, 2) = 0$.*

Proof. We focus on the second half. Figure (A1) corresponds to the following constraint:

$$\begin{aligned} \mu(3, 3) &= \nu(1, 3, 3) + \nu(2, 3, 3) + \nu(3, 3, 3) \\ &\leq \sum_a (\nu(1, 3, a) + \nu(2, 3, a)) + \sum_s \nu(3, s, 3) \\ &= (\mu(1, 1) + \mu(1, 2))p_{13} + (\mu(2, 1) + \mu(2, 2))p_{23} + (\mu(3, 1) + \mu(3, 2))p_{33} + \lambda_3 p_{33} \\ &= (\lambda_1 p_{13} + \lambda_2 p_{23} + \lambda_3 p_{33}) - (\mu(1, 3)p_{13} + \mu(2, 3)p_{23} + \mu(3, 3)p_{33}) + \lambda_3 p_{33} \\ &= \lambda_3(1 + p_{33}) - (\mu(1, 3)p_{13} + \mu(2, 3)p_{23} + \mu(3, 3)p_{33}) \\ &= \mu(1, 3)(1 + p_{33} - p_{13}) + \mu(2, 3)(1 + p_{23} - p_{13}) + \mu(3, 3). \end{aligned}$$

This can be an equality if and only if $\mu(1, 3) = \mu(2, 3) = 0$, which also implies $\mu(3, 1) = \mu(3, 2) = 0$. Likewise Figure (A2) corresponds to:

$$\begin{aligned}
\mu(3, 3) &= \nu(1, 3, 3) + \nu(2, 3, 3) + \nu(3, 3, 3) \\
&\leq \sum_s (\nu(1, s, 3) + \nu(2, s, 3)) + \sum_a \nu(3, 3, a) \\
&= \lambda_1 p_{13} + \lambda_2 p_{23} + (\mu(1, 3)p_{13} + \mu(2, 3)p_{23} + \mu(3, 3)p_{33}) \\
&= \lambda_3(1 - p_{33}) + (\mu(1, 3)p_{13} + \mu(2, 3)p_{23} + \mu(3, 3)p_{33}) \\
&= \mu(1, 3)(1 - p_{33} + p_{13}) + \mu(2, 3)(1 - p_{33} + p_{23}) + \mu(3, 3).
\end{aligned}$$

Again equality can only hold when $\mu(1, 3) = \mu(2, 3) = \mu(3, 1) = \mu(3, 2) = 0$. ■

Definition 12 *We call the constraints corresponding to (A1) to (A6) “type-I constraints centered at state 3.”*

Next we use this definition to investigate when two type-I constraints bind at μ^{tt} and μ :

Claim 8 *Suppose no four type-0 constraints bind at μ . Then among the type-I constraints centered at state 3, only two pairs (A3) and (A4) or (A5) and (A6) can simultaneously bind at μ . Moreover, when either pair binds simultaneously, the only type-0 constraint that can bind is $\mu(3, 1) = 0$ or $\mu(3, 2) = 0$.*

Proof. By the previous claim, we can ignore (A1) and (A2). If (A3) holds, then (A5) or (A6) cannot hold because that would imply $\mu(3, 1) = \mu(3, 2) = 0$. Thus (A3) is only compatible with (A4); similarly (A5) is only compatible with (A6).

For the second half, by symmetry we only need to consider (A3) and (A4) occurring simultaneously. From those figures we have:

$$\begin{aligned}
\mu(1, 3) + \mu(2, 3) &= \nu(1, 1, 3) + \nu(1, 2, 3) \\
&= \sum_s \nu(1, s, 3) - \sum_a \nu(1, 3, a) \\
&= \mu(2, 1)(p_{13} - p_{23}) + \mu(3, 1)(p_{13} - p_{33}).
\end{aligned}$$

It also holds that:

$$\begin{aligned}
\mu(1, 3) + \mu(2, 3) &= \mu(3, 1) + \mu(3, 2) \\
&= \nu(2, 3, 1) + \nu(2, 3, 2) \\
&= \sum_a \nu(2, 3, a) - \sum_s \nu(2, s, 3) \\
&= \mu(1, 2)(p_{13} - p_{23}) + \mu(3, 2)(p_{33} - p_{23}).
\end{aligned}$$

Since $\mu(1, 3) + \mu(2, 3) = \mu(3, 1) + \mu(3, 2)$, we deduce that $\mu(2, 1), \mu(1, 2)$ cannot be zero unless $\mu(1, 3) = \mu(2, 3) = 0$. Furthermore we have $\sum_a \nu(3, 3, a) = \nu(3, 3, 3) = \sum_s \nu(3, s, 3)$, thus

$$\mu(1, 3)(p_{13} - p_{33}) + \mu(2, 3)(p_{23} - p_{33}) = 0.$$

This suggests $\mu(1, 3)$ or $\mu(2, 3)$ cannot be zero unless they are both zero.³⁹ ■

Our next result handles the situation when two binding type-I constraints have different centers:

Claim 9 *Suppose $p_{ij} \leq \beta = \frac{\sqrt{5}-1}{2}$ and μ does not bind four type-0 constraints. If a type-I constraint centered at state 2 and another type-I constraint centered at state 3 both bind at μ^{tt} and μ , then the associated $\nu(a^{-1}, s, a)$ has one of the following two configurations:*

$$\nu = \begin{pmatrix} + & 0 & + \\ + & + & + \\ 0 & 0 & + \end{pmatrix} \begin{pmatrix} + & + & 0 \\ 0 & + & 0 \\ + & + & + \end{pmatrix} \begin{pmatrix} + & + & 0 \\ 0 & + & 0 \\ + & + & + \end{pmatrix} \quad (\text{A7})$$

$$\nu = \begin{pmatrix} + & + & 0 \\ 0 & + & 0 \\ + & + & + \end{pmatrix} \begin{pmatrix} + & 0 & + \\ + & + & + \\ 0 & 0 & + \end{pmatrix} \begin{pmatrix} + & 0 & + \\ + & + & + \\ 0 & 0 & + \end{pmatrix} \quad (\text{A8})$$

Moreover, when that happens $\mu(2, 3)$ and $\mu(3, 2)$ are the only type-0 constraints that can bind.

Proof. The four possible type-1 constraints centered at state 3 are shown in (A3) to (A6). These

³⁹We are ruling out the non-generic cases where $p_{13} = p_{33}$ and/or $p_{23} = p_{33}$. In those cases $\mu(1, 3)$ and/or $\mu(2, 3)$ might be zero, but such a type-0 constraint will not be independent from the two type-I constraints.

constraints bind precisely when:

$$\text{either } \mu(1, 3) + \mu(2, 3) = \pm (\mu(1, 2)(p_{23} - p_{13}) + \mu(3, 2)(p_{23} - p_{33})), \quad (47)$$

$$\text{or } \mu(1, 3) + \mu(2, 3) = \pm (\mu(2, 1)(p_{13} - p_{23}) + \mu(3, 1)(p_{13} - p_{33})). \quad (48)$$

Switching the states 2 and 3, we obtain that the four type-I constraints centered at state 2 bind when:

$$\text{either } \mu(1, 2) + \mu(3, 2) = \pm (\mu(1, 3)(p_{32} - p_{12}) + \mu(2, 3)(p_{32} - p_{22})), \quad (49)$$

$$\text{or } \mu(1, 2) + \mu(3, 2) = \pm (\mu(3, 1)(p_{12} - p_{32}) + \mu(2, 1)(p_{12} - p_{22})). \quad (50)$$

Ignoring the plus-minus signs, there are $2 \times 2 = 4$ possible combinations:

1. Suppose (47) and (49) hold at μ . The sum of the two equations gives $\mu(1, 2) = \mu(1, 3) = \mu(2, 3) = \mu(3, 2) = 0$, so $\mu = \mu^{tt}$.
2. Suppose (48) and (49) hold at μ . Consider (48) + $\beta \cdot$ (49):

$$\begin{aligned} & \mu(1, 3) + \mu(2, 3) + \mu(1, 2)\beta + \mu(3, 2)\beta \\ &= \pm [\mu(2, 1)(p_{13} - p_{23}) + \mu(3, 1)(p_{13} - p_{33})] \pm [\mu(1, 3)(p_{32} - p_{12})\beta + \mu(2, 3)(p_{32} - p_{22})\beta] \\ &\leq \mu(2, 1)\beta + \mu(3, 1)\beta + \mu(1, 3)\beta^2 + \mu(2, 3)\beta^2 \\ &= \mu(1, 2)\beta + \mu(1, 3)(\beta + \beta^2) + \mu(2, 3)\beta^2. \end{aligned}$$

Since $\beta + \beta^2 = 1$, this cannot hold unless $\mu(2, 1) = \mu(3, 1) = \mu(1, 3) = \mu(2, 3) = 0$.

3. Suppose (47) and (50) hold at μ . This is symmetric to the preceding case.
4. Finally suppose (48) and (50) hold at μ . By considering the sum and noting that $|p_{13} - p_{23} + p_{12} - p_{22}| < 1$, we deduce exactly one of the two equations takes the plus sign. This leads to two possible configurations for ν , which are exactly (A7) and (A8).

When (48) and (50) hold, $\mu(1, 2) = 0$ would imply $\mu(1, 3) = \mu(2, 1) + \mu(3, 1)$, contradicting (48). Similarly we can rule out $\mu(1, 3), \mu(2, 1), \mu(3, 1)$ being zero. This completes the proof. ■

To conclude the analysis of type-1 constraints, we show there cannot be 3 of them binding:

Claim 10 *Suppose $p_{ij} \leq \beta = \frac{\sqrt{5}-1}{2}$ and μ does not bind four type-0 constraints. There are at most 2 different type-I constraints that bind at μ^{tt} and μ .*

Proof. Suppose for contradiction that we have 3 type-I constraints binding. If any two of them have the same center, we can without loss assume the center is state 3. By Claim 8, the remaining binding constraint cannot be centered at state 3. But from Claim 9, it cannot be centered at state 2 (or 1) either, because any constraint centered at state 2 is compatible with at most one constraint centered at state 3. Either way we have reached a contradiction.

It remains to consider three type-I constraints with distinct centers. By Claim 9, we can assume ν has configuration (A7). Applying Claim 9 to states 1 and 3 instead, we obtain either $\nu(3, 3, 1) = \nu(3, 3, 2) = 0$ or $\nu(1, 1, 3) = \nu(1, 2, 3) = 0$. Together with (A7), the former implies another type-I constraint centered at state 3, contradiction to our assumption. The latter implies $\mu(1, 3) = \mu(2, 3) = 0$, again a contradiction. ■

G.2.3 Type-II constraints never bind

As discussed in Appendix D.2, a binding type-II constraint corresponds to 9 entries in $\nu(a^{-1}, s, a)$ to be zero. 8 of these entries form a $2 \times 2 \times 2$ “sub-cube”, see for example (25). Such a sub-cube involves two different values of s and two different values of a , which must overlap because $|S| = 3$. This means that a binding type-II constraint forces ν to be zero somewhere on the diagonal. We have thus shown:

Claim 11 *No type-II constraint binds at μ^{tt} .*

G.2.4 Type-III constraints never bind

While a type-III constraint also corresponds to 9 entries in $\nu(a^{-1}, s, a)$ to be zero, it differs from a type-II constraint in that exactly 3 entries are zero on each “face” of the $3 \times 3 \times 3$ cube. (26) is an example, but that will not bind at μ^{tt} because some diagonal entries are forced to be zero. With the goal of finding 9 off-diagonal entries to be zero, relatively straightforward enumeration using equation (23) yields:

Claim 12 *Up to relabelling of states, a type-III constraint binds at μ^{tt} if and only if the corre-*

sponding ν has one of the following configurations:

$$\nu = \begin{pmatrix} + & 0 & 0 \\ + & + & 0 \\ + & + & + \end{pmatrix} \begin{pmatrix} + & + & + \\ 0 & + & 0 \\ 0 & + & + \end{pmatrix} \begin{pmatrix} + & 0 & + \\ + & + & + \\ 0 & 0 & + \end{pmatrix} \quad (\text{A9})$$

$$\nu = \begin{pmatrix} + & 0 & + \\ + & + & + \\ 0 & 0 & + \end{pmatrix} \begin{pmatrix} + & 0 & 0 \\ + & + & 0 \\ + & + & + \end{pmatrix} \begin{pmatrix} + & + & + \\ 0 & + & 0 \\ 0 & + & + \end{pmatrix} \quad (\text{A10})$$

$$\nu = \begin{pmatrix} + & + & + \\ 0 & + & 0 \\ 0 & + & + \end{pmatrix} \begin{pmatrix} + & 0 & + \\ + & + & + \\ 0 & 0 & + \end{pmatrix} \begin{pmatrix} + & 0 & 0 \\ + & + & 0 \\ + & + & + \end{pmatrix} \quad (\text{A11})$$

$$\nu = \begin{pmatrix} + & 0 & 0 \\ + & + & 0 \\ + & + & + \end{pmatrix} \begin{pmatrix} + & 0 & + \\ + & + & + \\ 0 & 0 & + \end{pmatrix} \begin{pmatrix} + & + & + \\ 0 & + & 0 \\ 0 & + & + \end{pmatrix} \quad (\text{A12})$$

$$\nu = \begin{pmatrix} + & 0 & + \\ + & + & + \\ 0 & 0 & + \end{pmatrix} \begin{pmatrix} + & + & + \\ 0 & + & 0 \\ 0 & + & + \end{pmatrix} \begin{pmatrix} + & 0 & 0 \\ + & + & 0 \\ + & + & + \end{pmatrix} \quad (\text{A13})$$

$$\nu = \begin{pmatrix} + & + & + \\ 0 & + & 0 \\ 0 & + & + \end{pmatrix} \begin{pmatrix} + & 0 & 0 \\ + & + & 0 \\ + & + & + \end{pmatrix} \begin{pmatrix} + & 0 & + \\ + & + & + \\ 0 & 0 & + \end{pmatrix} \quad (\text{A14})$$

The first three can be ruled out relatively easily:

Claim 13 Suppose $\mu \neq \mu^{tt}$, then (A9) - (A11) does not occur at μ .

Proof. Suppose we have (A9). The position of zeros imply:

$$\begin{aligned} & \sum_a \nu(1, 1, a) + \sum_{a^{-1}} \nu(a^{-1}, 2, 3) + \sum_s \nu(2, s, 1) + \sum_s \nu(3, s, 1) + \sum_s \nu(3, s, 2) \\ &= \sum_{a^{-1}} \nu(a^{-1}, 1, 1) + \sum_a \nu(3, 2, a). \end{aligned}$$

This is equivalent to

$$\begin{aligned} & (\mu(1, 1)p_{11} + \mu(2, 1)p_{21} + \mu(3, 1)p_{31}) + \mu(2, 3) + \lambda_2 p_{21} + \lambda_3 p_{31} + \lambda_3 p_{32} \\ & = \mu(1, 1) + (\mu(1, 3)p_{12} + \mu(2, 3)p_{22} + \mu(3, 3)p_{32}). \end{aligned}$$

Using $\mu(1, 1) = \lambda_1 - \mu(2, 1) - \mu(3, 1)$ and $\mu(3, 3) = \lambda_3 - \mu(1, 3) - \mu(2, 3)$, we can simplify the above equation to:

$$\mu(2, 1)(1 - p_{11} + p_{21}) + \mu(3, 1)(1 - p_{11} + p_{31}) + \mu(2, 3)(1 - p_{22} + p_{32}) = \mu(1, 3)(p_{12} - p_{32}).$$

Since $\mu(2, 1) + \mu(3, 1) = \mu(1, 3) + \mu(1, 2)$ and $1 - p_{11} > p_{12}$, the L.H.S. is strictly larger than the R.H.S. unless $\mu(2, 1) = \mu(3, 1) = \mu(2, 3) = 0$. Thus $\mu = \mu^{tt}$ as desired.

For (A10), we can similarly obtain:

$$\mu(2, 1)(1 - p_{22} + p_{12}) + \mu(3, 1)(1 - p_{32} + p_{12}) + \mu(2, 3) = \mu(1, 2)(p_{21} - p_{11}) + \mu(3, 2)(p_{21} - p_{31}).$$

Since $\mu(2, 1) + \mu(2, 3) = \mu(1, 2) + \mu(3, 2)$ and $1 - p_{22} > p_{21}$, we again deduce $\mu = \mu^{tt}$.

For (A11), similar computation yields:

$$\mu(1, 2)(1 - p_{12} + p_{22}) + \mu(1, 3)(1 - p_{31} + p_{11}) + \mu(2, 3)(1 - p_{31} + p_{21}) = \mu(3, 2)(p_{32} - p_{22}).$$

Using $\mu(1, 3) + \mu(2, 3) = \mu(3, 1) + \mu(3, 2)$ and $1 - p_{31} > p_{32}$, we conclude $\mu = \mu^{tt}$. ■

It turns out that the remaining 3 type-III constraints can also be ruled about given the assumptions on p :

Claim 14 *Suppose $p_{ij} \leq \beta = \frac{\sqrt{5}-1}{2}$, and $\mu \neq \mu^{tt}$. Then (A12)–(A14) cannot occur at μ . Hence no type-III constraint binds at any $\mu \neq \mu^{tt}$.*

Proof. If (A12) occurs, we have:

$$\begin{aligned} & (\mu(1, 1)p_{11} + \mu(2, 1)p_{21} + \mu(3, 1)p_{31}) + \mu(2, 3) + \lambda_2 p_{21} + \lambda_2 p_{22} + \lambda_3 p_{31} \\ & = \mu(1, 1) + (\mu(1, 2)p_{12} + \mu(2, 2)p_{22} + \mu(3, 2)p_{32}). \end{aligned}$$

Substituting out the diagonal entries, we obtain:

$$\mu(2, 1)(1 - p_{11} + p_{21}) + \mu(3, 1)(1 - p_{11} + p_{31}) + \mu(2, 3) = \mu(1, 2)(p_{12} - p_{22}) + \mu(3, 2)(p_{32} - p_{22}).$$

Observe that $\mu(2, 1) + \mu(2, 3) = \mu(1, 2) + \mu(3, 2)$. If we can show $\min\{1 - p_{11} + p_{21}, 1\} > \max\{p_{12} - p_{22}, p_{32} - p_{22}\}$, then the above can only hold when $\mu = \mu^{tt}$. The only potential *caveat* is $1 - p_{11} + p_{21} > p_{32} - p_{22}$, or equivalently $p_{11} + p_{23} + p_{32} < 2$. So in fact it suffices to assume $p_{ij} < \frac{2}{3}$.

For (A13) and (A14), we simply note that they can be obtained from (A12) by permuting the states $1 \rightarrow 2 \rightarrow 3 \rightarrow 1$. This allows us to omit additional computation.⁴⁰ ■

G.2.5 Roadmap for the rest of the proof

Given the series of Claims 6 to 14, we now see that the three dual inequality constraints that bind at μ^{tt} and μ must be of type-0 or type-I, and they cannot all be type-I. There are a few possibilities:

1. In the next subsection we consider the case where μ satisfies three type-0 constraints.
2. After that we study the case with two type-0 constraints and one type-1 constraint. Up to relabelling of states, we could assume $\mu(1, 2) = 0$, and either $\mu(2, 1) = 0$ or $\mu(3, 1) = 0$. These are separately treated in two subsections.
3. Lastly we handle the case with one type-0 constraint and two type-1 constraints.

In each of these cases, we will directly construct τ and verify the conditions (40) to (46). This will allow us to prove $\mu \in \overline{\mathcal{C}_*}$ as intended.

G.3 Three type-0 constraints

In this subsection we assume that $\mu(s, a) = 0$ for at least three off-diagonal pairs (s, a) . There are two sub-cases: either two of the pairs lie on the same row or column, or the three pairs all have distinct s and a . In the first sub-case, we can without loss assume $\mu(1, 2) = \mu(1, 3) = \mu(2, 1) = \mu(3, 1) = 0$. By comparison, the second sub-case corresponds (ignoring permutation) to $\mu(1, 2) = \mu(2, 3) = \mu(3, 1) = 0$. We deal with these two sub-cases in order.

⁴⁰Note however that (A9)–(A11) are fixed points under this permutation. Thus the computations in the proof of the previous claim were not redundant.

G.3.1 $\mu(1, 2) = \mu(1, 3) = \mu(2, 1) = \mu(3, 1) = 0$

Since $\mu \in \mathcal{M}_1$, there is some $\nu^*(a^{-1}, s, a) \geq 0$ that solves the linear system defining \mathcal{M}_1 . We have $\nu^*(2, 1, 2) = \nu^*(2, 1, 3) = 0$, so:

$$\nu(2, 1, 1) = \sum_a \nu^*(2, 1, a) = \mu(1, 2)p_{11} + \mu(2, 2)p_{21} + \mu(3, 2)p_{31} = \mu(2, 2)p_{21} + \mu(3, 2)p_{31}.$$

Also because $\nu(2, 2, 1) = \nu(2, 3, 1) = 0$, we have:

$$\nu(2, 1, 1) = \sum_s \nu(2, s, 1) = \lambda_2 p_{21} = \mu(2, 2)p_{21} + \mu(3, 2)p_{21}.$$

Comparing these two equations for $\nu(2, 1, 1)$, we deduce $\mu(3, 2)(p_{31} - p_{21}) = 0$. As $\mu(3, 2) \neq 0$ (otherwise $\mu = \mu^{tt}$), we conclude $p_{31} = p_{21}$ just as in the proof of Theorem 1.

We now construct $\tau(s^{-1}, a^{-1}, s, a)$. Note that conditions (40) to (43) pin down the value of τ when $s^{-1} = a^{-1}$, so we will ignore them from now on. Moreover we may assume $\mu(s^{-1}, a^{-1}) > 0$, since otherwise τ can be arbitrarily specified. Thus in the current situation, we only need to find τ for $(s^{-1}, a^{-1}) = (2, 3), (3, 2)$. We verify that the following tables of $\tau(s, a | s^{-1}, a^{-1})$ work:

p_{21}	0	0	p_{21}	0	0
0	?	$\frac{\nu^*(1,2,3)+\nu^*(2,2,3)}{\mu(3,2)}$	0	?	$\frac{\nu^*(3,2,3)}{\mu(3,2)}$
0	$\frac{\nu^*(1,3,2)+\nu^*(2,3,2)}{\mu(3,2)}$?	0	$\frac{\nu^*(3,3,2)}{\mu(3,2)}$?
$s^{-1}=3, a^{-1}=2$			$s^{-1}=2, a^{-1}=3$		

(51)

We leave some entries with question marks to signify that their values can be solved from required row and column sums. But since (for example) $\tau(s = 2, a = 2 | s^{-1} = 3, a^{-1} = 2)$ can be solved from either the row sum or the column sum, we need to check that the same solution arises. This boils down to the following claim:

Claim 15 *For any solution ν^* to the linear system defining \mathcal{M}_1 , it holds that:*

$$\frac{\nu^*(1, 2, 3) + \nu^*(2, 2, 3)}{\mu(3, 2)} - \frac{\nu^*(1, 3, 2) + \nu^*(2, 3, 2)}{\mu(3, 2)} = p_{32} - p_{22}.$$

Proof. It suffices to show that $\nu^*(1, 2, 3) = \nu^*(1, 3, 2)$, while $\nu^*(2, 2, 3) - \nu^*(2, 3, 2) = \mu(3, 2)(p_{32} - p_{22})$.

p_{22}). Since $\nu^*(1, 2, 1) = \nu^*(1, 1, 2) = 0$, we have:

$$\nu^*(1, 2, 3) - \nu^*(1, 3, 2) = \sum_a \nu^*(1, 2, a) - \sum_s \nu^*(1, s, 2) = \mu(1, 1)p_{12} + \mu(2, 1)p_{22} + \mu(3, 1)p_{32} - \lambda_1 p_{12} = 0.$$

The last equality follows from $\mu(2, 1) = \mu(3, 1) = 0$. Similar calculation gives:

$$\begin{aligned} \nu^*(2, 2, 3) - \nu^*(2, 3, 2) &= \sum_a \nu^*(2, 2, a) - \sum_s \nu^*(2, s, 2) \\ &= \mu(1, 2)p_{12} + \mu(2, 2)p_{22} + \mu(3, 2)p_{32} - \lambda_2 p_{22} = \mu(3, 2)(p_{32} - p_{22}), \end{aligned}$$

which is exactly as claimed. ■

From this claim, row and column sums in the matrix $\tau(s, a | s^{-1} = 3, a^{-1} = 2)$ can be satisfied. Likewise we can have correct row and column sums in the second matrix in (51). Furthermore we see that condition (41) is satisfied for $(s, a) = (2, 3), (3, 2)$, as $\nu^*(1, 2, 3) + \nu^*(2, 2, 3) + \nu^*(3, 2, 3) = \mu(2, 3) = \mu(3, 2)$. Since it is trivially satisfied when $s = 1$ or $a = 1$, it is then satisfied everywhere. Finally τ is positive off the diagonal. By induction it's easy to check that condition (46) is satisfied for the ν defined by (44) and (45). This proves $\mu \in \overline{\mathcal{C}_*}$.

It is important to note that the ν induced by τ is generally not the same as the ν^* we started with, which could be any solution to the linear system. This is inevitable because when there are multiple $\nu^*(a^{-1}, s, a)$ that solve the linear system for \mathcal{M}_1 , often times some of them fail to extend to a solution to the linear system for \mathcal{M}_2 . We will see a recurring theme of choosing τ appropriately so that the resulting ν is positive.⁴¹ Henceforth in this appendix, we reserve ν^* to mean a solution to the linear system defining \mathcal{M}_1 , and ν to mean the distribution recursively defined by equations (44) and (45).

G.3.2 $\mu(1, 2) = \mu(2, 3) = \mu(3, 1) = 0$

Here we only need to find τ for $(s^{-1}, a^{-1}) = (2, 1), (3, 2), (1, 3)$. Take any solution ν^* to the linear system defining \mathcal{M}_1 , we try the following matrices for $\tau(s, a | s^{-1}, a^{-1})$ (note that

⁴¹This stands in contrast to the proof of $\mathcal{M}_1 = \mathcal{M}_k$, where τ is uniquely determined. The flexibility here comes in part from working with rays instead of vertices, the latter being constrained by one more dual inequalities.

$\mu(2, 1) = \mu(3, 2) = \mu(1, 3)$):

$$\begin{array}{c}
\begin{array}{|c|c|c|}
\hline
? & 0 & \frac{\nu^*(1,1,3)}{\mu(2,1)} \\
\hline
\frac{\nu^*(1,2,1)}{\mu(2,1)} & ? & 0 \\
\hline
0 & \frac{\nu^*(1,3,2)}{\mu(2,1)} & ? \\
\hline
\end{array} \\
s^{-1}=2, a^{-1}=1
\end{array}
\quad
\begin{array}{c}
\begin{array}{|c|c|c|}
\hline
? & 0 & \frac{\nu^*(2,1,3)}{\mu(3,2)} \\
\hline
\frac{\nu^*(2,2,1)}{\mu(3,2)} & ? & 0 \\
\hline
0 & \frac{\nu^*(2,3,2)}{\mu(3,2)} & ? \\
\hline
\end{array} \\
s^{-1}=3, a^{-1}=2
\end{array}
\quad
\begin{array}{c}
\begin{array}{|c|c|c|}
\hline
? & 0 & \frac{\nu^*(3,1,3)}{\mu(1,3)} \\
\hline
\frac{\nu^*(3,2,1)}{\mu(1,3)} & ? & 0 \\
\hline
0 & \frac{\nu^*(3,3,2)}{\mu(1,3)} & ? \\
\hline
\end{array} \\
s^{-1}=1, a^{-1}=3
\end{array}
\end{array} \quad (52)$$

As before, we should verify that the question marks can be filled out to have correct row and column sums simultaneously. Take as example the first matrix $\tau(s, a | s^{-1} = 2, a^{-1} = 1)$. The top-left question mark requires that $p_{21} - \frac{\nu^*(1,1,3)}{\mu(2,1)} = p_{11} - \frac{\nu^*(1,2,1)}{\mu(2,1)}$, equivalently $\nu^*(1, 1, 3) - \nu^*(1, 2, 1) = \mu(2, 1)(p_{21} - p_{11})$. Since $\nu^*(1, 1, 2) = \nu^*(1, 3, 1) = 0$, we have

$$\begin{aligned}
\nu^*(1, 1, 3) - \nu^*(1, 2, 1) &= \sum_a \nu^*(1, 1, a) - \sum_s \nu^*(1, s, 1) \\
&= \mu(1, 1)p_{11} + \mu(2, 1)p_{21} + \mu(3, 1)p_{31} - \lambda_1 p_{11} = \mu(2, 1)(p_{21} - p_{11}).
\end{aligned}$$

This is exactly as claimed, and the other question marks are treated analogously. Verifying condition (41) is straightforward, and the resulting ν is positive off the diagonal as τ is. Therefore all conditions are met and $\mu \in \overline{\mathcal{C}_*}$. In contrast to the previous sub-case, here we do have $\nu = \nu^*$.

G.4 Two type-0 constraints: $\mu(1, 2) = \mu(2, 1) = 0$.

The last dual constraint that binds at μ^{tt} and μ must be of type-I. If this type-I constraint is centered at state 3, then from the proof of Claim 9 we know that equation (47) or (48) must hold. But neither is consistent with $\mu(1, 2) = \mu(2, 1) = 0$ and $\mu(1, 3) + \mu(2, 3) = \mu(3, 1) + \mu(3, 2)$. Thus this type-I constraint must be centered at state 1 or 2. By symmetry we assume it is centered at state 1. From Claim 7 (switching states 1 and 3), we will then have 4 sub-cases for the zeros in ν^* :

$$\nu^* = \begin{pmatrix} + & 0 & + \\ 0 & + & + \\ 0 & + & + \end{pmatrix} \begin{pmatrix} + & 0 & + \\ 0 & + & + \\ 0 & + & + \end{pmatrix} \begin{pmatrix} + & 0 & 0 \\ 0 & + & + \\ + & + & + \end{pmatrix} \quad (B1)$$

$$\nu^* = \begin{pmatrix} + & 0 & 0 \\ 0 & + & + \\ + & + & + \end{pmatrix} \begin{pmatrix} + & 0 & 0 \\ 0 & + & + \\ + & + & + \end{pmatrix} \begin{pmatrix} + & 0 & + \\ 0 & + & + \\ 0 & + & + \end{pmatrix} \quad (B2)$$

$$\nu^* = \begin{pmatrix} + & 0 & 0 \\ 0 & + & + \\ + & + & + \end{pmatrix} \begin{pmatrix} + & 0 & + \\ 0 & + & + \\ 0 & + & + \end{pmatrix} \begin{pmatrix} + & 0 & 0 \\ 0 & + & + \\ + & + & + \end{pmatrix} \quad (\text{B3})$$

$$\nu^* = \begin{pmatrix} + & 0 & + \\ 0 & + & + \\ 0 & + & + \end{pmatrix} \begin{pmatrix} + & 0 & 0 \\ 0 & + & + \\ + & + & + \end{pmatrix} \begin{pmatrix} + & 0 & + \\ 0 & + & + \\ 0 & + & + \end{pmatrix} \quad (\text{B4})$$

We will treat these possibilities in turn.

G.4.1 Sub-case (B1)

The following computations are routine:

$$\nu^*(1, 1, 3) = \sum_a \nu^*(1, 1, a) - \sum_s \nu^*(1, s, 1) = \mu(3, 1)(p_{31} - p_{11}).$$

$$\nu^*(2, 1, 3) = \sum_a \nu^*(2, 1, a) - \sum_s \nu^*(2, s, 1) = \mu(3, 2)(p_{31} - p_{21}).$$

$$\nu^*(3, 3, 1) = \sum_s \nu^*(3, s, 1) - \sum_a \nu^*(3, 1, a) = \mu(1, 3)(p_{31} - p_{11}) + \mu(2, 3)(p_{31} - p_{21}).$$

From these we deduce $p_{31} \geq p_{11}, p_{21}$. We partially determine $\tau(s, a|s^{-1}, a^{-1})$ in the following way (ignoring those (s^{-1}, a^{-1}) with $s^{-1} = a^{-1}$ or $\mu = 0$):

$$\begin{array}{c} \begin{array}{|c|c|c|} \hline p_{11} & 0 & p_{31} - p_{11} \\ \hline 0 & ? & ? \\ \hline 0 & ? & ? \\ \hline \end{array} \\ s^{-1}=3, a^{-1}=1 \end{array} \quad \begin{array}{c} \begin{array}{|c|c|c|} \hline p_{21} & 0 & p_{31} - p_{21} \\ \hline 0 & ? & ? \\ \hline 0 & ? & ? \\ \hline \end{array} \\ s^{-1}=3, a^{-1}=2 \end{array} \quad \begin{array}{c} \begin{array}{|c|c|c|} \hline p_{11} & 0 & 0 \\ \hline 0 & ? & ? \\ \hline p_{31} - p_{11} & ? & ? \\ \hline \end{array} \\ s^{-1}=1, a^{-1}=3 \end{array} \quad \begin{array}{c} \begin{array}{|c|c|c|} \hline p_{21} & 0 & 0 \\ \hline 0 & ? & ? \\ \hline p_{31} - p_{21} & ? & ? \\ \hline \end{array} \\ s^{-1}=2, a^{-1}=3 \end{array} \quad (53)$$

As before, we follow the general rule to set $\tau(s, a|s^{-1}, a^{-1}) = 0$ when $\nu^*(a^{-1}, s, a) = 0$. So far we have the desired row and column sums in these matrices, and $\sum_{s^{-1}, a^{-1}} \tau(s^{-1}, a^{-1}, s, a) = \mu(s, a)$ holds when either $s = 1$ or $a = 1$. It remains to fill out the question marks to satisfy these conditions globally.

For easier reference, we denote the four matrices $\tau(\cdot, \cdot|s^{-1}, a^{-1})$ appearing in (53) by X, Y, Z, U , in the natural (left-to-right) order. Ignoring condition (41) for a moment, let us focus on con-

ditions (42) and (43). They impose row and column sums in each of the matrices. Because the first rows and columns are good, we simply need to worry about entries on the second and third row and column. More importantly, taking the matrix X as example, we can leave $X(2, 2)$ and $X(3, 3)$ undetermined so long as $X(2, 3) - X(3, 2) = p_{32} - p_{12}$. Their positivity do not matter because we only require ν to be positive off the diagonal. This analysis suggests the following choices to ensure the off-diagonal entries of X, Y, Z, W be positive (where x^+ denotes $\max\{x, 0\}$):

$$\begin{aligned}
X(2, 3) &= (p_{32} - p_{12})^+; & X(3, 2) &= X(2, 3) - (p_{32} - p_{12}). \\
Y(2, 3) &= (p_{32} - p_{22})^+; & Y(3, 2) &= Y(2, 3) - (p_{32} - p_{22}). \\
Z(2, 3) &= (p_{12} - p_{32})^+; & Z(3, 2) &= Z(2, 3) - (p_{12} - p_{32}). \\
U(2, 3) &= (p_{22} - p_{32})^+; & U(3, 2) &= U(2, 3) - (p_{22} - p_{32}).
\end{aligned} \tag{54}$$

We still need to check condition (41). Thanks to linear dependence among such equations, it suffices to prove $\sum_{s^{-1}, a^{-1}} \tau(s^{-1}, a^{-1}, s = 2, a = 3) = \mu(2, 3)$. Because adding an equal amount to $X(2, 3)$ and $X(3, 2)$ does not affect the other conditions, we only require the inequality version:

Claim 16 *Given $p_{ij} \leq \beta = \frac{\sqrt{5}-1}{2}$, it holds that:*

$$\mu(3, 1)X(2, 3) + \mu(3, 2)Y(2, 3) + \mu(1, 3)Z(2, 3) + \mu(2, 3)U(2, 3) \leq \mu(2, 3).$$

Proof. Since $\mu(1, 2) = \mu(2, 1) = 0$, we have $\mu(1, 3) = \mu(3, 1)$ and $\mu(2, 3) = \mu(3, 2)$. Plugging in (54), the desired claim becomes:

$$\mu(1, 3) \cdot |p_{32} - p_{12}| + \mu(2, 3) \cdot |p_{32} - p_{22}| \leq \mu(2, 3). \tag{55}$$

To prove (55), we discuss three possibilities:

1. $p_{12}, p_{22} \geq p_{32}$, or $p_{32} \geq p_{12}, p_{22}$. Here we claim that (55) follows from $\mu \in \mathcal{M}_1$. Indeed, routine computation yields:

$$\nu^*(1, 2, 3) - \nu^*(1, 3, 2) = \sum_a \nu^*(1, 2, a) - \sum_s \nu^*(1, s, 2) = \mu(3, 1)(p_{32} - p_{12}).$$

$$\nu^*(2, 2, 3) - \nu^*(2, 3, 2) = \sum_a \nu^*(2, 2, a) - \sum_s \nu^*(2, s, 2) = \mu(3, 2)(p_{32} - p_{22}).$$

$$\nu^*(3, 2, 3) - \nu^*(3, 3, 2) = \sum_a \nu^*(3, 2, a) - \sum_s \nu^*(3, s, 2) = \mu(1, 3)(p_{12} - p_{32}) + \mu(2, 3)(p_{22} - p_{32}).$$

From these we see that (55) follows immediately from the fact that $\nu^*(1, 2, 3) + \nu^*(2, 2, 3) + \nu^*(3, 2, 3) \leq \mu(2, 3)$.

For future reference, we call the above situation *constant signs* to signify that for any a^{-1} , the difference $\tau(s^{-1}, a^{-1}, 2, 3) - \tau(s^{-1}, a^{-1}, 3, 2)$ is either always positive or always negative, regardless of the values of $\mu(2, 3)$ and $\mu(3, 2)$.

2. $p_{12} > p_{32} > p_{22}$. We note that with configuration (B1),

$$\mu(1, 3) = \mu(3, 1) = \nu^*(3, 3, 1) = \mu(1, 3)(p_{31} - p_{11}) + \mu(2, 3)(p_{31} - p_{21}).$$

Thus we can deduce

$$\mu(1, 3)(1 - p_{31} + p_{11}) = \mu(2, 3)(p_{31} - p_{21}).$$

The desired inequality (55) is $\mu(1, 3)(p_{12} - p_{32}) \leq \mu(2, 3)(1 - p_{32} + p_{22})$, which then reduces to $(p_{31} - p_{21})(p_{12} - p_{32}) \leq (1 - p_{31} + p_{11})(1 - p_{32} + p_{22})$. This is because $p_{12} - p_{32} \leq p_{12}(1 - p_{32} + p_{22}) \leq \beta(1 - p_{32} + p_{22})$ and $\beta(p_{31} - p_{21}) \leq \beta^2 = 1 - \beta \leq 1 - p_{31} + p_{11}$.

3. $p_{22} > p_{32} > p_{12}$. In this case (55) becomes $\mu(1, 3)(p_{32} - p_{12}) \leq \mu(2, 3)(1 - p_{22} + p_{32})$, which further reduces to $(p_{31} - p_{21})(p_{32} - p_{12}) \leq (1 - p_{31} + p_{11})(1 - p_{22} + p_{32})$. In fact we have the stronger inequality $p_{31}p_{32} \leq (1 - p_{31})(1 - \beta + p_{32})$ whenever $p_{31}, p_{32} \leq \beta$. This is simply because $\beta p_{31} \leq 1 - p_{31}$ and $p_{32} \leq \beta(1 - \beta + p_{32})$.

We have thus proved (55) and the claim, which means $\mu \in \overline{\mathcal{C}_*}$. ■

G.4.2 Sub-case (B2)

Opposite to the previous sub-case, we have here $p_{31} \leq p_{11}, p_{21}$ by computing $\nu^*(1, 3, 1)$, $\nu^*(2, 3, 1)$ and $\nu^*(3, 1, 3)$. This suggests filling out τ partially as follows:

$$\begin{array}{c}
\begin{array}{|c|c|c|}
\hline
p_{31} & 0 & 0 \\
\hline
0 & ? & ? \\
\hline
p_{11} - p_{31} & ? & ? \\
\hline
\end{array} \\
s^{-1}=3, a^{-1}=1
\end{array}
\quad
\begin{array}{c}
\begin{array}{|c|c|c|}
\hline
p_{31} & 0 & 0 \\
\hline
0 & ? & ? \\
\hline
p_{21} - p_{31} & ? & ? \\
\hline
\end{array} \\
s^{-1}=3, a^{-1}=2
\end{array}
\quad
\begin{array}{c}
\begin{array}{|c|c|c|}
\hline
p_{31} & 0 & p_{11} - p_{31} \\
\hline
0 & ? & ? \\
\hline
0 & ? & ? \\
\hline
\end{array} \\
s^{-1}=1, a^{-1}=3
\end{array}
\quad
\begin{array}{c}
\begin{array}{|c|c|c|}
\hline
p_{31} & 0 & p_{21} - p_{31} \\
\hline
0 & ? & ? \\
\hline
0 & ? & ? \\
\hline
\end{array} \\
s^{-1}=2, a^{-1}=3
\end{array}
\end{array} \tag{56}$$

Again we assign $X(2, 3), X(3, 2)$ and etc., according to (54), so that row and column sums can be satisfied and τ is positive off the diagonal. As before, we need to check the validity of Claim 16, or equivalently of the inequality (55).

Again we distinguish three possibilities:

1. $p_{12}, p_{22} \geq p_{32}$, or $p_{32} \geq p_{12}, p_{22}$. We have “constant signs” as discussed in the previous sub-case, and (55) follows from $\mu \in \mathcal{M}_1$.
2. $p_{12} > p_{32} > p_{22}$. From configuration (B2) we obtain $\mu(1, 3) = \nu^*(3, 1, 3) = \mu(1, 3)(p_{11} - p_{31}) + \mu(2, 3)(p_{21} - p_{31})$. Thus:

$$\mu(1, 3)(1 - p_{11} + p_{31}) = \mu(2, 3)(p_{21} - p_{31}).$$

Using this relation, inequality (55) reduces to $(p_{21} - p_{31})(p_{12} - p_{32}) \leq (1 - p_{11} + p_{31})(1 - p_{32} + p_{22})$. This follows from $p_{12} - p_{32} \leq p_{12}(1 - p_{32} + p_{22}) \leq \beta(1 - p_{32} + p_{22})$, and $\beta(p_{21} - p_{31}) \leq \beta^2 \leq 1 - p_{11} + p_{31}$.

3. $p_{22} > p_{32} > p_{12}$. Using again the relation between $\mu(1, 3)$ and $\mu(2, 3)$, we reduce (55) to $(p_{21} - p_{31})(p_{32} - p_{12}) \leq (1 - p_{11} + p_{31})(1 - p_{22} + p_{32})$. This follows from $\beta p_{32} \leq (1 - \beta)(1 - \beta + p_{32})$ whenever $p_{32} \leq \beta$.

Thus the vertex μ in this sub-case also belongs to $\overline{\mathcal{C}}_*$.

G.4.3 Sub-case (B3)

We compute that:

$$\nu^*(1, 3, 1) = \mu(3, 1)(p_{11} - p_{31}).$$

$$\nu^*(2, 1, 3) = \mu(3, 2)(p_{31} - p_{21}).$$

$$\nu^*(3, 3, 1) = \mu(1, 3)(p_{31} - p_{11}) + \mu(2, 3)(p_{31} - p_{21}).$$

Thus $p_{11} \geq p_{31} \geq p_{21}$ and we determine τ partially as follows:

p_{31}	0	0	p_{21}	0	$p_{31} - p_{21}$	p_{11}	0	0	p_{21}	0	0
0	?	?	0	?	?	0	?	?	0	?	?
$p_{11} - p_{31}$?	?	0	?	?	$p_{31} - p_{11}$?	?	$p_{31} - p_{21}$?	?
$s^{-1}=3, a^{-1}=1$			$s^{-1}=3, a^{-1}=2$			$s^{-1}=1, a^{-1}=3$			$s^{-1}=2, a^{-1}=3$		

(57)

The major difference from before is that τ necessarily has an entry that is negative: $Z(3, 1) = p_{31} - p_{11} \leq 0$. This complicates the analysis because we want the resulting ν to be positive. Fortunately, we can still apply induction to show $\nu \geq 0$ so long as for every s^{-1}, a^{-1} :

$$\tau(s = 2, a = 3 | s^{-1}, a^{-1})(p_{31} - p_{21}) \geq \tau(s = 1, a = 3 | s^{-1}, a^{-1})(p_{11} - p_{31}).$$

This suggests we should additionally require $Y(2, 3) \geq p_{11} - p_{31}$. We thus change the specification of Y in (54) and impose that:

$$X(2, 3) = (p_{32} - p_{12})^+; \quad Y(2, 3) = \max\{p_{32} - p_{22}, p_{11} - p_{31}\}; \quad Z(2, 3) = (p_{12} - p_{32})^+; \quad U(2, 3) = (p_{22} - p_{32})^+. \quad (58)$$

As before, we shall verify the crucial Claim 16, which in this case becomes:

$$\mu(1, 3) \cdot |p_{32} - p_{12}| + \mu(2, 3) \cdot (\max\{p_{32} - p_{22}, p_{11} - p_{31}\} + (p_{22} - p_{32})^+) \leq \mu(2, 3). \quad (59)$$

There are four possibilities:

1. $p_{12}, p_{22} \geq p_{32}$. Observe from configuration (B3) that:

$$\mu(1, 3) = \nu^*(2, 1, 3) = \mu(3, 2)(p_{31} - p_{21}) = \mu(2, 3)(p_{31} - p_{21}).$$

Thus (59) reduces to $\mu(1, 3)(p_{12} - p_{32}) \leq \mu(2, 3)(1 - p_{11} + p_{31} - p_{22} + p_{32}) = \mu(2, 3)(2 - p_{11} - p_{22} - p_{33})$. Using the above relation between $\mu(1, 3)$ and $\mu(2, 3)$, this is equivalent to

$$(p_{31} - p_{21})(p_{12} - p_{32}) \leq 2 - p_{11} - p_{22} - p_{33}.$$

This follows from $(1 - p_{33})(1 - p_{11}) \leq 2 - \beta - p_{11} - p_{33}$ whenever $p_{11}, p_{33} \leq \beta$.

2. $p_{32} \geq p_{12}, p_{22}$. Then (59) reduces to $\mu(1, 3)(p_{32} - p_{12}) + \mu(2, 3) \max\{p_{32} - p_{22}, p_{11} - p_{31}\} \leq \mu(2, 3)$, which further simplifies to $(p_{31} - p_{21})(p_{32} - p_{12}) \leq \min\{1 - p_{32} + p_{22}, 1 - p_{11} + p_{31}\}$. This holds because the L.H.S. is at most $\beta^2 = 1 - \beta$, while the R.H.S. is at least $1 - \beta$.
3. $p_{12} > p_{32} > p_{22}$. Similarly (59) reduces to $(p_{31} - p_{21})(p_{12} - p_{32}) \leq \min\{1 - p_{32} + p_{22}, 1 - p_{11} + p_{31}\}$, with the sign of $p_{12} - p_{32}$ switched from before. This inequality holds for the same reason as before.

4. $p_{22} > p_{32} > p_{12}$. Here (59) reduces to $(p_{31} - p_{21})(p_{32} - p_{12}) \leq 2 - p_{11} - p_{22} - p_{33}$. To prove this inequality, note that the L.H.S. is at most $\frac{(1-p_{33})^2}{4} \leq \frac{1-p_{33}}{4}$ by AM-GM inequality, while the R.H.S. is at least $2 - 2\beta - p_{33}$. We have $\frac{1-p_{33}}{4} < 2 - 2\beta - p_{33}$ whenever $p_{33} \leq \beta$.

We have thus proved inequality (59). This means we can find τ that satisfies conditions (40) to (43). Furthermore, by construction τ is positive off the diagonal except for $Z(3, 1)$, so that $\nu(a^{-k}, \dots, a^{-1}, s, a)$ is positive by induction whenever $(a^{-1}, s, a) \neq (3, 3, 1)$. If $(a^{-1}, s, a) = (3, 3, 1)$, we can use $\tau(s = 2, a = 3|s^{-1}, a^{-1})(p_{31} - p_{21}) \geq \tau(s = 1, a = 3|s^{-1}, a^{-1})(p_{11} - p_{31})$ and second-order induction to prove $\nu \geq 0$, just as we did in the proof of $\mathcal{M}_1 = \mathcal{M}_k$. Hence $\mu \in \overline{C}_*$ as desired.

G.4.4 Sub-case (B4)

Opposite to the previous sub-case, here we have $p_{21} \geq p_{31} \geq p_{11}$. We can partially determine the conditional τ as follows:

$$\begin{array}{c}
\begin{array}{|c|c|c|} \hline p_{11} & 0 & p_{31} - p_{11} \\ \hline 0 & ? & ? \\ \hline 0 & ? & ? \\ \hline \end{array} \\
s^{-1}=3, a^{-1}=1
\end{array}
\quad
\begin{array}{c}
\begin{array}{|c|c|c|} \hline p_{31} & 0 & 0 \\ \hline 0 & ? & ? \\ \hline p_{21} - p_{31} & ? & ? \\ \hline \end{array} \\
s^{-1}=3, a^{-1}=2
\end{array}
\quad
\begin{array}{c}
\begin{array}{|c|c|c|} \hline p_{31} & 0 & p_{11} - p_{31} \\ \hline 0 & ? & ? \\ \hline 0 & ? & ? \\ \hline \end{array} \\
s^{-1}=1, a^{-1}=3
\end{array}
\quad
\begin{array}{c}
\begin{array}{|c|c|c|} \hline p_{31} & 0 & p_{21} - p_{31} \\ \hline 0 & ? & ? \\ \hline 0 & ? & ? \\ \hline \end{array} \\
s^{-1}=2, a^{-1}=3
\end{array}
\tag{60}$$

These choices ensure that the row and column sums are correct for the first row and column. Just as in the previous sub-case, τ already has a negative entry: $Z(1, 3) = p_{11} - p_{31} \leq 0$. Despite this, we know that one way to ensure the positivity of ν is to have:

$$\tau(s = 2, a = 3|s^{-1}, a^{-1})(p_{21} - p_{31}) \geq \tau(s = 1, a = 3|s^{-1}, a^{-1})(p_{31} - p_{11}). \tag{61}$$

With a bit of foresight, we set $X(2, 3), Y(2, 3), U(2, 3)$ to satisfy the above constraint with equality and leave $Z(2, 3)$ to be determined later (partly since $Z(2, 3)$ could also be negative). Specifically we choose:

$$\begin{aligned}
X(2, 3) &= \max\{p_{32} - p_{12}, \frac{(p_{31} - p_{11})^2}{p_{21} - p_{31}}\}; & X(3, 2) &= X(2, 3) - (p_{32} - p_{12}). \\
Y(2, 3) &= (p_{32} - p_{22})^+; & Y(3, 2) &= Y(2, 3) - (p_{32} - p_{22}). \\
U(2, 3) &= \max\{p_{22} - p_{32}, p_{31} - p_{11}\}; & U(3, 2) &= U(2, 3) - (p_{22} - p_{32}).
\end{aligned}
\tag{62}$$

In each of the following four possibilities, we choose $Z(2, 3)$ carefully to satisfy Claim 16 and simultaneously ensure $\nu \geq 0$:

1. $p_{12}, p_{22} \geq p_{32}$. Here we let $Z(2, 3) = p_{12} - p_{32}$ and accordingly $Z(3, 2) = 0$. This way, τ has only one negative entry $Z(1, 3)$, and (61) holds for every s^{-1}, a^{-1} . Thus by induction ν is positive. It remains to check Claim 16, which by (62) reduces to:

$$\mu(1, 3) \cdot \left(\frac{(p_{31} - p_{11})^2}{p_{21} - p_{31}} + p_{12} - p_{32} \right) + \mu(2, 3) \cdot \max\{p_{22} - p_{32}, p_{31} - p_{11}\} \leq \mu(2, 3).$$

From configuration (B4) we have:

$$\mu(1, 3) = \mu(3, 1) = \nu^*(2, 3, 1) = \mu(3, 2)(p_{21} - p_{31}) = \mu(2, 3)(p_{21} - p_{31}).$$

Thus the desired inequality further reduces to $(p_{31} - p_{11})^2 + (p_{21} - p_{31})(p_{12} - p_{32}) + \max\{p_{22} - p_{32}, p_{31} - p_{11}\} \leq 1$. This holds because the L.H.S. is at most $p_{31}^2 + (\beta - p_{31})\beta + \beta \leq 1$.

2. $p_{32} \geq p_{12}, p_{22}$. Here we let $Z(2, 3) = -\frac{(p_{31} - p_{11})^2}{p_{21} - p_{31}}$ and $Z(3, 2) = Z(2, 3) - (p_{12} - p_{32})$. This ensures that (61) holds for every s^{-1}, a^{-1} . Thus by induction, $\nu(\dots, a^{-1} = 3, s = 1, a = 3) \geq 0$. Note however that $Z(2, 3)$ is now negative, raising the concern that $\nu(\dots, a^{-1} = 3, s = 2, a = 3)$ might be negative. We rule this out by observing that $Z(2, 3)/Z(1, 3) = U(2, 3)/U(1, 3) = \frac{p_{31} - p_{11}}{p_{21} - p_{31}}$. Thus $\nu(\dots, 3, 2, 3) = \frac{p_{31} - p_{11}}{p_{21} - p_{31}} \cdot \nu(\dots, 3, 1, 3)$, which is positive. Furthermore $Z(3, 2) = Z(2, 3) + (p_{32} - p_{12})$ could also be negative, making it possible that $\nu(\dots, a^{-1} = 3, s = 3, a = 2)$ be negative. This also does not happen because we can deduce from $Z(3, 2) \geq Z(2, 3)$ and $U(3, 2) \geq U(2, 3)$ that $\nu(\dots, 3, 3, 2) \geq \nu(\dots, 3, 2, 3)$.

Hence we do have $\nu \geq 0$. It remains to check Claim 16, which by (62) boils down to:

$$\mu(1, 3) \cdot \left(p_{32} - p_{12} - \frac{(p_{31} - p_{11})^2}{p_{21} - p_{31}} \right)^+ + \mu(2, 3) \cdot (p_{32} - p_{22} + p_{31} - p_{11}) \leq \mu(2, 3).$$

Using $\mu(1, 3) = (p_{21} - p_{31})\mu(2, 3)$, the above further simplifies to

$$\left((p_{32} - p_{12})(p_{21} - p_{31}) - (p_{31} - p_{11})^2 \right)^+ \leq p_{11} + p_{22} + p_{33}.$$

To prove this inequality, note that $p_{32} - p_{12} = p_{32} + p_{13} + p_{11} - p_{13} \leq 2\beta - 1 + p_{11}$. Also

$p_{21} - p_{31} \leq \beta - p_{31}$ and $p_{33} \geq 1 - \beta - p_{31}$. It thus suffices to show $(2\beta - 1 + p_{11})(\beta - p_{31}) \leq p_{11} + 1 - \beta - p_{31} + (p_{31} - p_{11})^2$. Collecting terms, this becomes $0 \leq p_{11}^2 + (1 - \beta - p_{31})p_{11} + (p_{31}^2 - (2 - 2\beta)p_{31} + 1 - 2\beta^2)$. If $p_{31} \leq 1 - \beta$, the first two summands are positive and the last summand is at least $1 - 2\beta^2 - (1 - \beta)^2 = 2\beta - 3\beta^2 > 0$, so we are done.

If $p_{31} > 1 - \beta$, instead of using $p_{33} \geq 1 - \beta - p_{31}$ we simply use $p_{33} \geq 0$. It then suffices to show $(2\beta - 1 + p_{11})(\beta - p_{31}) \leq p_{11} + (p_{31} - p_{11})^2$. This is equivalent to $0 \leq p_{31}^2 + (2\beta - 1 - p_{11})p_{31} + (p_{11}^2 + (1 - \beta)p_{11} - \beta(2\beta - 1))$. As a function of p_{31} , the derivative of the R.H.S. is $2p_{31} + 2\beta - 1 - p_{11} \geq 0$ when $p_{31} \geq 1 - \beta$. Thus to prove the last inequality we can assume $p_{31} = 1 - \beta$. But this then reduce to our previous analysis for $p_{31} \leq 1 - \beta$.

Therefore Claim 16 holds either way, and the τ we construct do satisfy all desired conditions.

3. $p_{12} > p_{32} > p_{22}$. From (62) we have $X(2, 3) = \frac{(p_{31} - p_{11})^2}{p_{21} - p_{31}}$, $Y(2, 3) = p_{32} - p_{22}$ and $U(2, 3) = p_{31} - p_{11}$. Let us take:

$$Z(2, 3) = \frac{-(p_{31} - p_{11})^2 + p_{11} + p_{22} + p_{33}}{p_{21} - p_{31}}.$$

Then Claim 16 reduces to:

$$\mu(2, 3)(p_{32} - p_{22} + p_{11} + p_{22} + p_{33} + p_{31} - p_{11}) \leq \mu(2, 3).$$

This is in fact an equality. It remains to check the positivity of ν . Observe that the only negative entries in τ are $Z(1, 3)$, $Z(2, 3)$ and $Z(3, 2) = Z(2, 3) - (p_{12} - p_{32})$. Thus by induction we only need to prove the positivity of $\nu(a^{-k}, \dots, a^{-1}, s, a)$ for $a^{-1} = 3$ and $(s, a) = (1, 3)$, $(2, 3)$ or $(3, 2)$.

Firstly, $\nu(a^{-k}, \dots, a^{-1} = 3, s = 1, a = 3) \geq 0$ follows from the usual argument that $\tau(s = 2, a = 3 | s^{-1}, a^{-1})(p_{21} - p_{31}) \geq \tau(s = 1, a = 3 | s^{-1}, a^{-1})(p_{31} - p_{11})$ for every s^{-1}, a^{-1} .

Secondly, we have $Z(2, 3) \geq \frac{p_{31} - p_{11}}{p_{21} - p_{31}} Z(1, 3)$ and $U(2, 3) \geq \frac{p_{31} - p_{11}}{p_{21} - p_{31}} U(1, 3)$. Thus:

$$\nu(a^{-k}, \dots, a^{-1} = 3, s = 2, a = 3) \geq \frac{p_{31} - p_{11}}{p_{21} - p_{31}} \cdot \nu(a^{-k}, \dots, a^{-1} = 3, s = 1, a = 3) \geq 0.$$

In fact, we have the more precise relation (which will be useful later):

$$\begin{aligned} & \nu(a^{-k}, \dots, a^{-1} = 3, s = 2, a = 3) \\ &= \frac{p_{31} - p_{11}}{p_{21} - p_{31}} \cdot \nu(a^{-k}, \dots, a^{-1} = 3, s = 1, a = 3) + \frac{p_{11} + p_{22} + p_{33}}{p_{21} - p_{31}} \cdot \nu(a^{-k}, \dots, a^{-2}, s^{-1} = 1, a^{-1} = 3). \end{aligned}$$

Lastly we verify $\nu(\dots, 3, 3, 2) \geq 0$. From the recursive relation (45) we have:

$$\begin{aligned} & \nu(a^{-k}, \dots, a^{-1} = 3, s = 3, a = 2) - \nu(a^{-k}, \dots, a^{-1} = 3, s = 2, a = 3) \\ &= (Z(3, 2) - Z(2, 3)) \cdot \nu(\dots, s^{-1} = 1, a^{-1} = 3) + (U(3, 2) - U(2, 3)) \cdot \nu(\dots, s^{-1} = 2, a^{-1} = 3) \\ &= (p_{32} - p_{22}) \cdot \nu(a^{-k}, \dots, a^{-2}, s^{-1} = 2, a^{-1} = 3) - (p_{12} - p_{32}) \cdot \nu(a^{-k}, \dots, a^{-2}, s^{-1} = 1, a^{-1} = 3) \end{aligned}$$

Combining the previous two equations, we obtain:

$$\begin{aligned} & \nu(a^{-k}, \dots, a^{-1} = 3, s = 3, a = 2) \\ & \geq (p_{32} - p_{22}) \cdot \nu(\dots, s^{-1} = 2, a^{-1} = 3) - \left(p_{12} - p_{32} - \frac{p_{11} + p_{22} + p_{33}}{p_{21} - p_{31}} \right) \cdot \nu(\dots, s^{-1} = 1, a^{-1} = 3). \end{aligned}$$

Since we already know $\nu(a^{-k}, \dots, a^{-2}, s^{-1} = 2, a^{-1} = 3) \geq \frac{p_{31} - p_{11}}{p_{21} - p_{31}} \cdot \nu(a^{-k}, \dots, a^{-2}, s^{-1} = 1, a^{-1} = 3)$. The above equation suggests $\nu(\dots, 3, 3, 2)$ is positive if:⁴²

$$(p_{32} - p_{22})(p_{31} - p_{11}) \geq (p_{12} - p_{32})(p_{21} - p_{31}) - p_{11} - p_{22} - p_{33}.$$

Writing $p_{11} + p_{22} + p_{33} = 1 - (p_{32} - p_{22}) - (p_{31} - p_{11})$, this inequality then reduces to $(1 - p_{32} + p_{12})(1 - p_{31} + p_{11}) \geq (p_{12} - p_{32})(p_{21} - p_{31})$, which obviously holds. Hence ν is indeed positive.

4. $p_{22} > p_{32} > p_{12}$. Here we choose $Z(2, 3) = 0$, so that τ only has one negative entry off the diagonal, $Z(1, 3)$. This way, $\nu \geq 0$ is not a problem. By (62), Claim 16 reduces to:

$$\mu(1, 3) \cdot \max\left\{p_{32} - p_{12}, \frac{(p_{31} - p_{11})^2}{p_{21} - p_{31}}\right\} + \mu(2, 3) \cdot \max\{p_{22} - p_{32}, p_{31} - p_{11}\} \leq \mu(2, 3).$$

Using $\mu(1, 3) = (p_{21} - p_{31})\mu(2, 3)$, the above follows easily from $\beta^2 + \beta = 1$.

⁴²If we had simply chosen $Z(2, 3) = \frac{-(p_{31} - p_{11})^2}{p_{21} - p_{31}}$ as before, we would have to deduce the positivity of ν from $(p_{32} - p_{22})(p_{31} - p_{11}) \geq (p_{12} - p_{32})(p_{21} - p_{31})$, which is not generally true.

Hence for all parameter values we have constructed τ to prove $\mu \in \overline{\mathcal{C}_*}$. This sub-case is completed, and so is the entire subsection. We now turn to a different situation with two type-0 constraints.

G.5 Two type-0 constraints: $\mu(1, 2) = \mu(3, 1) = 0$

Like in the previous subsection, there is a type-I constraint binding at both μ and μ^{tt} . Suppose this type-I constraint is centered at state 3. Then from the proof of Claim 9 we know that equation (47) or (48) must hold. But neither of them is consistent with $\mu(1, 2) = \mu(3, 1) = 0$ (which implies $\mu(3, 2) = \mu(1, 3) + \mu(2, 3) = \mu(2, 1) + \mu(2, 3)$). Similarly we cannot have (49) or (50), so that the type-I constraint cannot be centered at state 2, either. Thus the type-I constraint must be centered at state 1, and we again have four sub-cases to consider. The configurations of $\nu^*(a^{-1}, s, a)$ are shown below, analogous to (B1) to (B4) before:

$$\nu^* = \begin{pmatrix} + & 0 & + \\ 0 & + & + \\ 0 & + & + \end{pmatrix} \begin{pmatrix} + & 0 & + \\ 0 & + & + \\ 0 & + & + \end{pmatrix} \begin{pmatrix} + & 0 & 0 \\ + & + & + \\ 0 & + & + \end{pmatrix} \quad (\text{B5})$$

$$\nu^* = \begin{pmatrix} + & 0 & 0 \\ + & + & + \\ 0 & + & + \end{pmatrix} \begin{pmatrix} + & 0 & 0 \\ + & + & + \\ 0 & + & + \end{pmatrix} \begin{pmatrix} + & 0 & + \\ 0 & + & + \\ 0 & + & + \end{pmatrix} \quad (\text{B6})$$

$$\nu^* = \begin{pmatrix} + & 0 & 0 \\ + & + & + \\ 0 & + & + \end{pmatrix} \begin{pmatrix} + & 0 & + \\ 0 & + & + \\ 0 & + & + \end{pmatrix} \begin{pmatrix} + & 0 & 0 \\ + & + & + \\ 0 & + & + \end{pmatrix} \quad (\text{B7})$$

$$\nu^* = \begin{pmatrix} + & 0 & + \\ 0 & + & + \\ 0 & + & + \end{pmatrix} \begin{pmatrix} + & 0 & 0 \\ + & + & + \\ 0 & + & + \end{pmatrix} \begin{pmatrix} + & 0 & + \\ 0 & + & + \\ 0 & + & + \end{pmatrix} \quad (\text{B8})$$

G.5.1 Sub-case (B5)

Routine computation of $\nu^*(1, 1, 3)$, $\nu^*(2, 1, 3)$ and $\nu^*(3, 2, 1)$ shows $p_{31} \geq p_{21} \geq p_{11}$. We can partially determine $\tau(s, a|s^{-1}, a^{-1})$ as follows:

$$\begin{array}{c}
\begin{array}{|c|c|c|}
\hline
p_{11} & 0 & p_{21} - p_{11} \\
\hline
0 & ? & ? \\
\hline
0 & ? & ? \\
\hline
\end{array} \\
s^{-1}=2, a^{-1}=1
\end{array}
\quad
\begin{array}{c}
\begin{array}{|c|c|c|}
\hline
p_{21} & 0 & p_{31} - p_{21} \\
\hline
0 & ? & ? \\
\hline
0 & ? & ? \\
\hline
\end{array} \\
s^{-1}=3, a^{-1}=2
\end{array}
\quad
\begin{array}{c}
\begin{array}{|c|c|c|}
\hline
p_{11} & 0 & 0 \\
\hline
p_{31} - p_{11} & ? & ? \\
\hline
0 & ? & ? \\
\hline
\end{array} \\
s^{-1}=1, a^{-1}=3
\end{array}
\quad
\begin{array}{c}
\begin{array}{|c|c|c|}
\hline
p_{21} & 0 & 0 \\
\hline
p_{31} - p_{21} & ? & ? \\
\hline
0 & ? & ? \\
\hline
\end{array} \\
s^{-1}=2, a^{-1}=3
\end{array}
\quad (63)
\end{array}$$

This is similar to (53), the major difference being that the first matrix now corresponds to $s^{-1} = 2, a^{-1} = 1$. Let us still call the four matrices X, Y, Z, U . So far τ is all positive, and has correct row and column sums. To ensure these properties globally, we set:

$$\begin{aligned}
X(2, 3) &= (p_{22} - p_{12})^+; & X(3, 2) &= X(2, 3) - (p_{22} - p_{12}). \\
Y(2, 3) &= (p_{32} - p_{22})^+; & Y(3, 2) &= Y(2, 3) - (p_{32} - p_{22}). \\
Z(2, 3) &= (p_{33} - p_{13})^+; & Z(3, 2) &= Z(2, 3) - (p_{33} - p_{13}). \\
U(2, 3) &= (p_{33} - p_{23})^+; & U(3, 2) &= U(2, 3) - (p_{33} - p_{23}).
\end{aligned} \tag{64}$$

Compared with equation (54), the above choices are different in that we have not only transition probabilities on the second column, but those on the third column as well. With τ given by (63), we do not need to worry about the positivity of ν . It thus remains to verify the following analogue of Claim 16:

Claim 17 *Given $p_{ij} \leq \beta = \frac{\sqrt{5}-1}{2}$, it holds that:*

$$\mu(2, 1) \cdot X(2, 3) + \mu(3, 2) \cdot Y(2, 3) + \mu(1, 3) \cdot Z(2, 3) + \mu(2, 3) \cdot U(2, 3) \leq \mu(2, 3).$$

Proof. We discuss three possibilities:

1. $p_{13}, p_{23} \geq p_{33}$ or $p_{13}, p_{23} \leq p_{33}$. We have constant signs and the claim follows from $\mu \in \mathcal{M}_1$.
2. $p_{13} > p_{33} > p_{23}$. Then $Z(2, 3) = 0$. We also deduce from $p_{31} \geq p_{21}$ and $p_{33} > p_{23}$ that $p_{32} < p_{22}$, and so $Y(2, 3) = 0$. From configuration (B5), we obtain that $\mu(1, 3) = \mu(2, 1) = \nu^*(3, 2, 1) = \mu(1, 3)(p_{31} - p_{11}) + \mu(2, 3)(p_{31} - p_{21})$. Thus we have the following relation:

$$\mu(2, 1)(1 - p_{31} + p_{11}) = \mu(2, 3)(p_{31} - p_{21}).$$

Plugging these in, the desired claim reduces to $\mu(2, 1) \cdot (p_{22} - p_{12})^+ + \mu(2, 3) \cdot (p_{33} - p_{23}) \leq$

$\mu(2, 3)$, which further simplifies to $(1 - p_{31} + p_{11})(1 - p_{33} + p_{23}) \geq (p_{31} - p_{21})(p_{22} - p_{12})$. Since $p_{21} \geq p_{11}$ and $p_{22} - p_{12} = p_{22} + p_{11} + p_{13} - 1 \leq 2\beta - 1 + p_{11}$. It suffices to show $(1 - \beta + p_{11})(1 - \beta) \geq (\beta - p_{11})(2\beta - 1 + p_{11})$. Collecting terms, this last inequality becomes $1 - \beta - \beta^2 + p_{11}^2 \geq 0$, which obviously holds.

3. $p_{23} > p_{33} > p_{13}$. Then $U(2, 3) = 0$. We also deduce from $p_{21} \geq p_{11}$ and $p_{23} > p_{13}$ that $p_{22} < p_{12}$, and so $X(2, 3) = 0$. The claim reduces to $\mu(3, 2)(p_{32} - p_{22})^+ + \mu(1, 3)(p_{33} - p_{13}) \leq \mu(2, 3)$. If $p_{32} \leq p_{22}$, then as $\mu(1, 3) = \mu(2, 1)$ we need to show $(p_{31} - p_{21})(p_{33} - p_{13}) \leq 1 - p_{31} + p_{11}$. This follows from $\beta^2 \leq 1 - \beta$. If instead $p_{32} > p_{22}$, then as $\mu(3, 2) = \mu(2, 1) + \mu(2, 3)$ we need to show $(p_{31} - p_{21})(p_{32} - p_{22} + p_{33} - p_{13}) \leq (1 - p_{31} + p_{11})(1 - p_{32} + p_{22})$. This holds because $p_{31} - p_{21} \leq 1 - p_{32} + p_{22}$ and $p_{33} - p_{13} + p_{32} - p_{22} \leq 1 - p_{31} + p_{11}$.

We have thus proved Claim 17 for all parameter values. $\mu \in \overline{\mathcal{C}}_*$ as desired. ■

G.5.2 Sub-case (B6)

Here we have $p_{11} \geq p_{21} \geq p_{31}$. τ can be partially determined as follows:

$$\begin{array}{c}
\begin{array}{|c|c|c|}
\hline
p_{21} & 0 & 0 \\
\hline
p_{11} - p_{21} & ? & ? \\
\hline
0 & ? & ? \\
\hline
\end{array} \\
s^{-1}=2, a^{-1}=1
\end{array}
\quad
\begin{array}{c}
\begin{array}{|c|c|c|}
\hline
p_{31} & 0 & 0 \\
\hline
p_{21} - p_{31} & ? & ? \\
\hline
0 & ? & ? \\
\hline
\end{array} \\
s^{-1}=3, a^{-1}=2
\end{array}
\quad
\begin{array}{c}
\begin{array}{|c|c|c|}
\hline
p_{31} & 0 & p_{11} - p_{31} \\
\hline
0 & ? & ? \\
\hline
0 & ? & ? \\
\hline
\end{array} \\
s^{-1}=1, a^{-1}=3
\end{array}
\quad
\begin{array}{c}
\begin{array}{|c|c|c|}
\hline
p_{31} & 0 & p_{21} - p_{31} \\
\hline
0 & ? & ? \\
\hline
0 & ? & ? \\
\hline
\end{array} \\
s^{-1}=2, a^{-1}=3
\end{array}
\end{array} \tag{65}$$

The following choices ensure that τ has correct row and column sums and is positive off the diagonal:

$$\begin{aligned}
X(2, 3) &= (p_{13} - p_{23})^+; & X(3, 2) &= X(2, 3) - (p_{13} - p_{23}). \\
Y(2, 3) &= (p_{23} - p_{33})^+; & Y(3, 2) &= Y(2, 3) - (p_{23} - p_{33}). \\
Z(2, 3) &= (p_{12} - p_{32})^+; & Z(3, 2) &= Z(2, 3) - (p_{12} - p_{32}). \\
U(2, 3) &= (p_{22} - p_{32})^+; & U(3, 2) &= U(2, 3) - (p_{22} - p_{32}).
\end{aligned} \tag{66}$$

It remains to check Claim 17. Again there are three possibilities:

1. $p_{12}, p_{22} \geq p_{32}$ or $p_{12}, p_{22} \leq p_{32}$. We have constant signs, so no more proof is needed.
2. $p_{12} > p_{32} > p_{22}$. Then $U(2, 3) = 0$. We also deduce from $p_{11} \geq p_{21}$ and $p_{12} > p_{22}$ that $p_{13} < p_{23}$, and so $X(2, 3) = 0$. Observe from configuration (B6) that $\mu(1, 3) = \nu^*(3, 1, 3) =$

$\mu(1, 3)(p_{11} - p_{31}) + \mu(2, 3)(p_{21} - p_{31})$. Thus:

$$\mu(1, 3)(1 - p_{11} + p_{31}) = \mu(2, 3)(p_{21} - p_{31}).$$

Claim 17 reduces to $\mu(3, 2)(p_{23} - p_{33})^+ + \mu(1, 3)(p_{12} - p_{32}) \leq \mu(2, 3)$. If $p_{23} < p_{33}$, this becomes $(p_{21} - p_{31})(p_{12} - p_{32}) \leq 1 - p_{11} + p_{31}$, which follows from $\beta^2 \leq 1 - \beta$. If instead $p_{23} > p_{33}$, then as $\mu(3, 2) = \mu(1, 3) + \mu(2, 3)$ we need to show $(p_{21} - p_{31})(p_{23} - p_{33} + p_{12} - p_{32}) \leq (1 - p_{11} + p_{31})(1 - p_{23} + p_{33})$. This holds because $p_{21} - p_{31} \leq 1 - p_{23} + p_{33}$ and $p_{12} - p_{32} + p_{23} - p_{33} = p_{12} + p_{23} + p_{31} - 1 \leq 1 - p_{11} + p_{31}$.

3. $p_{22} > p_{32} > p_{12}$, Then $Z(2, 3) = 0$. We also deduce from $p_{21} \geq p_{31}$ and $p_{22} > p_{32}$ that $p_{23} < p_{33}$, and so $Y(2, 3) = 0$. Claim 17 reduces to $\mu(2, 1)(p_{13} - p_{23})^+ + \mu(2, 3)(p_{22} - p_{32}) \leq \mu(2, 3)$. By $\mu(2, 1) = \mu(1, 3)$ and the preceding relation between $\mu(1, 3)$ and $\mu(2, 3)$, we need to show $(p_{21} - p_{31})(p_{13} - p_{23}) \leq (1 - p_{11} + p_{31})(1 - p_{22} + p_{32})$. This follows from $p_{21} - p_{31} \leq 1 - p_{22} + p_{32}$ and $p_{13} - p_{23} \leq 1 - p_{11} + p_{31}$.

We have thus resolved this sub-case as well.

G.5.3 Sub-case (B7)

We compute that:

$$\nu^*(1, 2, 1) = \mu(2, 1)(p_{11} - p_{21}).$$

$$\nu^*(2, 1, 3) = \mu(3, 2)(p_{31} - p_{21}).$$

$$\nu^*(3, 2, 1) = \mu(1, 3)(p_{31} - p_{11}) + \mu(2, 3)(p_{31} - p_{21}).$$

Thus $p_{11}, p_{31} \geq p_{21}$. We set the conditional τ partially as follows:

$$\begin{array}{c}
\begin{array}{|c|c|c|}
p_{21}	0	0
$p_{11} - p_{21}$?	?
0	?	?
$s^{-1}=2, a^{-1}=1$		
\end{array}		
\quad		
\begin{array}{c}		
\begin{array}{	c	c
p_{21}	0	$p_{31} - p_{21}$
0	?	?
0	?	?
$s^{-1}=3, a^{-1}=2$		
\end{array}		
\quad		
\begin{array}{c}		
\begin{array}{	c	c
p_{11}	0	0
$p_{31} - p_{11}$?	?
0	?	?
$s^{-1}=1, a^{-1}=3$		
\end{array}		
\quad		
\begin{array}{c}		
\begin{array}{	c	c
p_{21}	0	0
$p_{31} - p_{21}$?	?
0	?	?
 $s^{-1}=2, a^{-1}=3$
\end{array}
\tag{67}$$

The situation is reminiscent of sub-case (B3), but here $Z(2, 1) = p_{31} - p_{11}$ may or may not be positive. Let us first assume $p_{31} \geq p_{11}$, so this is not a concern. We then make the following

choices:

$$\begin{aligned}
X(2, 3) &= (p_{13} - p_{23})^+; & X(3, 2) &= X(2, 3) - (p_{13} - p_{23}). \\
Y(2, 3) &= (p_{32} - p_{22})^+; & Y(3, 2) &= Y(2, 3) - (p_{32} - p_{22}). \\
Z(2, 3) &= (p_{33} - p_{13})^+; & Z(3, 2) &= Z(2, 3) - (p_{33} - p_{13}). \\
U(2, 3) &= (p_{33} - p_{23})^+; & U(3, 2) &= U(2, 3) - (p_{33} - p_{23}).
\end{aligned} \tag{68}$$

Let us verify Claim 17 for various parameter values:

1. $p_{13}, p_{23} \geq p_{33}$ or $p_{13}, p_{23} \leq p_{33}$. We are done due to constant signs.
2. $p_{13} > p_{33} > p_{23}$. Then $Z(2, 3) = 0$. We also deduce from $p_{33} > p_{23}$ and $p_{31} \geq p_{21}$ that $p_{32} < p_{22}$, and so $Y(2, 3) = 0$. Claim 17 reduces to:

$$\mu(2, 1)(p_{13} - p_{23}) + \mu(2, 3)(p_{33} - p_{23}) \leq \mu(2, 3).$$

Now observe from configuration (B7) that $\mu(1, 3) = \nu^*(2, 1, 3) = \mu(3, 2)(p_{31} - p_{21}) = (\mu(1, 3) + \mu(2, 3))(p_{31} - p_{21})$. Thus:

$$\mu(1, 3)(1 - p_{31} + p_{21}) = \mu(2, 3)(p_{31} - p_{21}).$$

Using this relation, we only need to show $(p_{31} - p_{21})(p_{13} - p_{23}) \leq (1 - p_{31} + p_{21})(1 - p_{33} + p_{23})$. This holds because $p_{31} - p_{21} \leq 1 - p_{33} + p_{23}$ and $p_{13} - p_{23} = p_{13} + p_{21} + p_{22} - 1 \leq 1 - p_{31} + p_{21}$.

3. $p_{23} > p_{33} > p_{13}$. Then $X(2, 3) = U(2, 3) = 0$. Claim 17 reduces to $\mu(3, 2) \cdot (p_{32} - p_{22})^+ + \mu(1, 3)(p_{33} - p_{13}) \leq \mu(2, 3)$. If $p_{32} \leq p_{22}$, this further simplifies to $(p_{31} - p_{21})(p_{33} - p_{13}) \leq 1 - p_{31} + p_{21}$, which follows from $\beta^2 \leq 1 - \beta$. If instead $p_{32} > p_{22}$, we need to show $(p_{31} - p_{21})(p_{32} - p_{22} + p_{33} - p_{13}) \leq (1 - p_{31} + p_{21})(1 - p_{32} + p_{22})$. This holds because $p_{31} - p_{21} \leq 1 - p_{32} + p_{22}$ and $p_{32} - p_{22} + p_{33} - p_{13} \leq 1 - p_{31} + p_{21}$.

Thus when $p_{31} \geq p_{11}$, we can find τ to satisfy all the conditions for $\mu \in \overline{\mathcal{C}_*}$. Next we tackle the slightly more tricky situation when $p_{31} < p_{11}$. As usual, we are going to impose for every s^{-1}, a^{-1} :

$$\tau(s = 2, a = 3 | s^{-1}, a^{-1}) \geq \frac{p_{11} - p_{31}}{p_{31} - p_{21}} \cdot \tau(s = 1, a = 3 | s^{-1}, a^{-1}).$$

This is sufficient to ensure $\nu \geq 0$. Note that we only additionally need $Y(2, 3) \geq p_{11} - p_{31}$. Thus we modify the choices in (68) to set:

$$Y(2, 3) = \max\{p_{32} - p_{22}, p_{11} - p_{31}\}.$$

Below we verify that Claim 17 still holds under this modification. Without loss we assume $Y(2, 3) = p_{11} - p_{31}$; otherwise we can resort to previous analysis. Now we distinguish four possibilities:

- (a) $p_{13}, p_{23} \geq p_{33}$. Then from (68) $Z(2, 3) = U(2, 3) = 0$. Claim 17 reduces to $\mu(2, 1)(p_{13} - p_{23})^+ + \mu(3, 2)(p_{11} - p_{31}) \leq \mu(2, 3)$. If $p_{13} < p_{23}$, this further simplifies (using the above relation between $\mu(1, 3)$ and $\mu(2, 3)$) to $(p_{31} - p_{21})(p_{11} - p_{31}) \leq (1 - p_{31} + p_{21})(1 - p_{11} + p_{31})$. This is equivalent to $0 \leq 1 - p_{11} + p_{21}$, which obviously holds. If instead $p_{13} \geq p_{23}$, we need to show $(p_{31} - p_{21})(p_{13} - p_{23} + p_{11} - p_{31}) \leq (1 - p_{31} + p_{21})(1 - p_{11} + p_{31})$. This holds because $p_{31} - p_{21} \leq 1 - p_{11} + p_{31}$ and $p_{13} - p_{23} + p_{11} - p_{31} \leq 1 - p_{31} + p_{21}$.
- (b) $p_{13}, p_{23} \leq p_{33}$. If $p_{23} \geq p_{13}$, Claim 17 reduces to $\mu(3, 2)(p_{11} - p_{31}) + \mu(1, 3)(p_{33} - p_{13}) + \mu(2, 3)(p_{33} - p_{23}) \leq \mu(2, 3)$. This is equivalent to:

$$\mu(1, 3)(p_{11} - p_{31} + p_{33} - p_{13}) \leq \mu(2, 3)(1 - p_{11} + p_{31} - p_{33} + p_{23}).$$

The above further simplifies to $(p_{31} - p_{21})(p_{11} - p_{31} + p_{33} - p_{13}) \leq (1 - p_{31} + p_{21})(1 - p_{11} + p_{31} - p_{33} + p_{23})$, which is just $(p_{31} - p_{21})(p_{23} - p_{13}) \leq 1 - p_{11} + p_{21} - p_{33} + p_{23} = 2 - p_{11} - p_{22} - p_{33}$. This holds because the L.H.S. is at most $(1 - p_{33})(1 - p_{22}) = 1 - p_{22} - p_{33} + p_{22}p_{33}$, which is no more than the R.H.S. when $p_{ij} \leq \beta$.⁴³

If instead $p_{23} < p_{13}$, then Claim 17 reduces to

$$\mu(1, 3)(p_{11} - p_{31} + p_{33} - p_{23}) \leq \mu(2, 3)(1 - p_{11} + p_{31} - p_{33} + p_{23}).$$

This ends up being $0 \leq 2 - p_{11} - p_{22} - p_{33}$, which obviously holds.

⁴³The inequality here is tight when $p_{11} = p_{22} = p_{33} = \beta$ and $p_{13} = p_{21} = p_{32} = 0$. As we show later, once the persistence p_{ii} exceeds β , the μ considered here will belong to \mathcal{C}_1 but not \mathcal{C}_2 . In this sense our choices of τ are almost optimal.

- (c) $p_{13} > p_{33} > p_{23}$. From (68) we have $X(2, 3) = p_{13} - p_{23}$, $Y(2, 3) = p_{11} - p_{31}$, $Z(2, 3) = 0$ and $U(2, 3) = p_{33} - p_{23}$. Then Claim 17 becomes

$$\mu(2, 1)(p_{13} - p_{23}) + \mu(3, 2)(p_{11} - p_{31}) + \mu(2, 3)(p_{33} - p_{23}) \leq \mu(2, 3).$$

Writing $\mu(2, 3) = \mu(3, 2) - \mu(2, 1)$ and using $\mu(2, 1) = \mu(3, 2)(p_{31} - p_{21})$, we need to show $(p_{31} - p_{21})(1 + p_{13} - p_{33}) \leq 1 - p_{11} + p_{31} - p_{33} + p_{23}$. This simplifies to $(p_{31} - p_{21})(p_{13} - p_{33}) \leq 1 - p_{11} + p_{21} - p_{33} + p_{23} = 2 - p_{11} - p_{22} - p_{33}$, which holds because the L.H.S. is at most $(1 - p_{33})(1 - p_{11})$.

- (d) $p_{23} > p_{33} > p_{13}$. We have from (68) $X(2, 3) = 0$, $Y(2, 3) = p_{11} - p_{31}$, $Z(2, 3) = p_{33} - p_{13}$ and $U(2, 3) = 0$. The desired claim becomes

$$\mu(3, 2)(p_{11} - p_{31}) + \mu(1, 3)(p_{33} - p_{13}) \leq \mu(2, 3) = \mu(3, 2) - \mu(1, 3).$$

Since $\mu(1, 3) = (p_{31} - p_{21})\mu(3, 2)$, we need to show $(p_{31} - p_{21})(1 + p_{33} - p_{13}) \leq 1 - p_{11} + p_{31}$. This is equivalent to $(p_{31} - p_{21})(p_{33} - p_{13}) \leq 1 - p_{11} + p_{21}$, which follows from $\beta^2 \leq 1 - \beta$.

We have completed the analysis of this sub-case.

G.5.4 Sub-case (B8)

Here we have $\nu^*(1, 1, 3) = \mu(2, 1)(p_{21} - p_{11})$, $\nu^*(1, 2, 1) = \mu(3, 2)(p_{21} - p_{31})$ and $\nu^*(3, 1, 3) = \mu(1, 3)(p_{11} - p_{31}) + \mu(2, 3)(p_{21} - p_{31})$. Thus $p_{21} \geq p_{11}, p_{31}$. The conditional τ can be partially determined as follows:

$$\begin{array}{c}
\begin{array}{|c|c|c|}
\hline
p_{11} & 0 & p_{21} - p_{11} \\
\hline
0 & ? & ? \\
\hline
0 & ? & ? \\
\hline
\end{array} \\
s^{-1}=2, a^{-1}=1
\end{array}
\quad
\begin{array}{c}
\begin{array}{|c|c|c|}
\hline
p_{31} & 0 & 0 \\
\hline
p_{21} - p_{31} & ? & ? \\
\hline
0 & ? & ? \\
\hline
\end{array} \\
s^{-1}=3, a^{-1}=2
\end{array}
\quad
\begin{array}{c}
\begin{array}{|c|c|c|}
\hline
p_{31} & 0 & p_{11} - p_{31} \\
\hline
0 & ? & ? \\
\hline
0 & ? & ? \\
\hline
\end{array} \\
s^{-1}=1, a^{-1}=3
\end{array}
\quad
\begin{array}{c}
\begin{array}{|c|c|c|}
\hline
p_{31} & 0 & p_{21} - p_{31} \\
\hline
0 & ? & ? \\
\hline
0 & ? & ? \\
\hline
\end{array} \\
s^{-1}=2, a^{-1}=3
\end{array}
\end{array} \tag{69}$$

Note that $Z(1, 3) = p_{11} - p_{31}$ may or may not be negative. Like in the previous sub-case, we first suppose $p_{11} \geq p_{31}$ so this will not be a concern. We make τ positive off the diagonal by

choosing:

$$\begin{aligned}
X(2, 3) &= (p_{22} - p_{12})^+; & X(3, 2) &= X(2, 3) - (p_{22} - p_{12}). \\
Y(2, 3) &= (p_{23} - p_{33})^+; & Y(3, 2) &= Y(2, 3) - (p_{23} - p_{33}). \\
Z(2, 3) &= (p_{12} - p_{32})^+; & Z(3, 2) &= Z(2, 3) - (p_{12} - p_{32}). \\
U(2, 3) &= (p_{22} - p_{32})^+; & U(3, 2) &= U(2, 3) - (p_{22} - p_{32}).
\end{aligned} \tag{70}$$

We shall verify Claim 17 by considering three possibilities:

1. $p_{12}, p_{22} \geq p_{32}$ or $p_{12}, p_{22} \leq p_{32}$. No more proof is needed because we have constant signs.
2. $p_{12} > p_{32} > p_{22}$. Then $X(2, 3) = U(2, 3) = 0$. Claim 17 reduces to:

$$\mu(3, 2)(p_{23} - p_{33})^+ + \mu(1, 3)(p_{12} - p_{32}) \leq \mu(2, 3) = \mu(3, 2) - \mu(1, 3).$$

Now observe from configuration (B8) that:

$$\mu(1, 3) = \mu(2, 1) = \nu^*(2, 2, 1) = \mu(3, 2)(p_{21} - p_{31}).$$

Using this relation, we just need to show $(p_{21} - p_{31})(1 + p_{12} - p_{32}) + (p_{23} - p_{33})^+ \leq 1$. If $p_{23} \leq p_{33}$, this follows from $\beta(1 + \beta) \leq 1$. If instead $p_{23} > p_{33}$, the last inequality is equivalent to $(p_{21} - p_{31})(p_{12} - p_{32}) \leq 1 - p_{21} + p_{31} - p_{23} + p_{33} = 1 + p_{22} - p_{32}$. This holds again because $\beta^2 \leq 1 - \beta$.

3. $p_{22} > p_{32} > p_{12}$. Here we have $Z(2, 3) = 0$. We also deduce from $p_{21} \geq p_{31}$ and $p_{22} > p_{32}$ that $p_{23} < p_{33}$, and so $Y(2, 3) = 0$. Claim 17 becomes $\mu(2, 1)(p_{22} - p_{12}) + \mu(2, 3)(p_{22} - p_{32}) \leq \mu(2, 3)$. Using $\mu(2, 1) = \mu(1, 3)$, $\mu(2, 3) = \mu(3, 2) - \mu(1, 3)$ and $\mu(1, 3) = \mu(3, 2)(p_{21} - p_{31})$, we need to show $(p_{21} - p_{31})(p_{22} - p_{12}) \leq (1 - p_{21} + p_{31})(1 - p_{22} + p_{32})$. This holds because $p_{21} - p_{31} \leq 1 - p_{22} + p_{32}$ and $p_{22} - p_{12} \leq 1 - p_{21} + p_{31}$.

Thus when $p_{11} \geq p_{31}$, we can find τ to satisfy the desired conditions. We now turn to the more difficult situation when $p_{21} \geq p_{31} > p_{11}$. As $\mu(1, 3) = \mu(2, 1) = \nu^*(2, 2, 1) = \mu(3, 2)(p_{21} - p_{31}) > 0$, we actually have the strict inequality $p_{21} > p_{31}$. To ensure the positivity of ν , we routinely impose:

$$\tau(s = 2, a = 3 | s^{-1}, a^{-1}) \geq \frac{p_{31} - p_{11}}{p_{21} - p_{31}} \cdot \tau(s = 1, a = 3 | s^{-1}, a^{-1}).$$

Like in sub-case (B4), we choose X, Y, U to exactly satisfy the above constraint, and leave $Z(2, 3)$ to be determined later.

$$\begin{aligned} X(2, 3) &= \max\left\{p_{22} - p_{12}, \frac{(p_{21} - p_{11})(p_{31} - p_{11})}{p_{21} - p_{31}}\right\}, \\ Y(2, 3) &= (p_{23} - p_{33})^+. \\ U(2, 3) &= \max\{p_{22} - p_{32}, p_{31} - p_{11}\}. \end{aligned} \tag{71}$$

In each of the following four possibilities, we choose $Z(2, 3)$ to satisfy Claim 17 and simultaneously ensure $\nu \geq 0$:

- (a) $p_{12}, p_{22} \geq p_{32}$. Here we set $Z(2, 3) = p_{12} - p_{32}$, so that τ only has one negative entry off the diagonal. By construction, ν is positive despite this negative entry. It remains to check Claim 17. As $p_{21} > p_{31}$ and $p_{22} \geq p_{32}$, we have $p_{23} < p_{33}$. Thus $Y(2, 3) = 0$ and Claim 17 reduces to:

$$\begin{aligned} &\mu(1, 3) \cdot \max\left\{p_{22} - p_{12}, \frac{(p_{21} - p_{11})(p_{31} - p_{11})}{p_{21} - p_{31}}\right\} \\ &+ \mu(1, 3) \cdot (p_{12} - p_{32}) + \mu(2, 3) \cdot \max\{p_{22} - p_{32}, p_{31} - p_{11}\} \leq \mu(2, 3). \end{aligned} \tag{72}$$

Suppose both maximum operators select the first term, then we need to show $\mu(1, 3)(p_{22} - p_{32}) + \mu(2, 3)(p_{22} - p_{32}) \leq \mu(2, 3)$. As $\mu(2, 3) = \mu(3, 2) - \mu(1, 3)$ and $\mu(1, 3) = \mu(3, 2)(p_{21} - p_{31})$, the last inequality is equivalent to $p_{22} - p_{32} \leq 1 - p_{21} + p_{31}$, which obviously holds.

Suppose the first max is $p_{22} - p_{12}$ while the second max is $p_{31} - p_{11}$, then we need to show $\mu(1, 3)(p_{22} - p_{32}) + \mu(2, 3)(p_{31} - p_{11}) \leq \mu(2, 3)$. This is equivalent to $(p_{21} - p_{31})(p_{22} - p_{32}) \leq (1 - p_{21} + p_{31})(1 - p_{31} + p_{11})$, or simply $(p_{21} - p_{31})(p_{11} + p_{22} + p_{33}) \leq 1 - p_{31} + p_{11}$. This follows from $(p_{21} - p_{31})p_{11} \leq p_{11}$ and $(p_{21} - p_{31})(p_{22} + p_{33}) \leq p_{21}(1 - p_{31})(p_{22} + p_{33}) < 1 - p_{31}$.

Suppose the first max in (72) is achieved by the second term while the second max is $p_{22} - p_{32}$. Then we need to show $\mu(3, 2)(p_{21} - p_{11})(p_{31} - p_{11}) + \mu(1, 3)(p_{12} - p_{32}) + \mu(2, 3)(p_{22} - p_{32}) \leq \mu(2, 3)$. This is equivalent to $(p_{21} - p_{11})(p_{31} - p_{11}) + (p_{21} - p_{31})(p_{12} - p_{32}) \leq (1 - p_{21} + p_{31})(1 - p_{22} + p_{32})$. Some simple manipulation further reduces the inequality to $(p_{21} - p_{11})(p_{31} - p_{11}) + (p_{21} - p_{31})(p_{12} - p_{22}) \leq 1 - p_{21} + p_{31} - p_{22} + p_{32} = 1 + p_{23} - p_{33}$. This inequality holds because it is linear in p_{31} on the interval $[p_{11}, p_{21}]$, and it holds at the extreme points because $\beta^2 \leq 1 - \beta$.

Lastly suppose both max in (72) are achieved by the second argument. Then we need to show $\mu(3, 2)(p_{21} - p_{11})(p_{31} - p_{11}) + \mu(1, 3)(p_{12} - p_{32}) + \mu(2, 3)(p_{31} - p_{11}) \leq \mu(2, 3)$. This is equivalent to $(p_{21} - p_{11})(p_{31} - p_{11}) + (p_{21} - p_{31})(p_{12} - p_{32}) \leq (1 - p_{21} + p_{31})(1 - p_{31} + p_{11})$. Collecting terms, this becomes $(p_{21} - p_{31})(p_{11} + p_{12} + p_{33}) \leq 1 - (p_{31} - p_{11})(1 + p_{21} - p_{11})$, or $(p_{21} - p_{31})(p_{12} + p_{33}) \leq 1 - p_{31}(1 + p_{21}) + p_{11}(1 + 2p_{31} - p_{11})$. In fact the stronger inequality $(p_{21} - p_{31})(p_{12} + p_{33}) \leq 1 - p_{31}(1 + p_{21})$ holds, because it is linear in p_{31} and holds at the $p_{31} = 0$ and $p_{31} = p_{21}$.

- (b) $p_{12}, p_{22} \leq p_{32}$. Here we let $Z(2, 3) = \frac{p_{31} - p_{11}}{p_{21} - p_{31}} Z(1, 3) = -\frac{(p_{31} - p_{11})^2}{p_{21} - p_{31}}$. Like in sub-case (B4), this choice ensures that $\nu(a^{-k}, \dots, a^{-1} = 3, s = 1, a = 3) \geq 0$. We also have:

$$\nu(a^{-k}, \dots, a^{-1} = 3, s = 2, a = 3) = \frac{p_{31} - p_{11}}{p_{21} - p_{31}} \cdot \nu(a^{-k}, \dots, a^{-1} = 3, s = 1, a = 3) \geq 0.$$

$$\nu(a^{-k}, \dots, a^{-1} = 3, s = 3, a = 2) \geq \nu(a^{-k}, \dots, a^{-1} = 3, s = 2, a = 3) \geq 0.$$

Thus the possibly negative entries in τ , $Z(1, 3)$, $Z(2, 3)$ and $Z(3, 2)$, do not affect the positivity of ν . It remains to check Claim 17, which by (71) becomes:

$$\begin{aligned} & \mu(2, 1) \cdot \max\left\{p_{22} - p_{12}, \frac{(p_{21} - p_{11})(p_{31} - p_{11})}{p_{21} - p_{31}}\right\} + \mu(3, 2) \cdot (p_{23} - p_{33})^+ \\ & - \mu(1, 3) \cdot \frac{(p_{31} - p_{11})^2}{p_{21} - p_{31}} + \mu(2, 3) \cdot (p_{31} - p_{11}) \leq \mu(2, 3). \end{aligned}$$

Because $\mu(2, 3) = \mu(3, 2) - \mu(1, 3)$ and $\mu(1, 3) = (p_{21} - p_{31})\mu(3, 2)$, the above simplifies to:

$$\max\{(p_{21} - p_{31})(p_{22} - p_{12}), (p_{21} - p_{11})(p_{31} - p_{11})\} + (p_{23} - p_{33})^+ + (p_{31} - p_{11})(1 - p_{21} + p_{11}) \leq 1 - p_{21} + p_{31}.$$

Suppose the max is achieved by the first term, then the L.H.S. is at most $p_{22} + p_{23} + p_{31}$, which is just the R.H.S. Suppose the max is achieved by the second term, then we need to show $(p_{23} - p_{33})^+ \leq 1 - p_{21} + p_{11}$, which obviously holds.

- (c) $p_{12} > p_{32} > p_{22}$. Here we know from (71) that $X(2, 3) = \frac{(p_{21} - p_{11})(p_{31} - p_{11})}{p_{21} - p_{31}}$, $Y(2, 3) = (p_{23} - p_{33})^+$ and $U(2, 3) = p_{31} - p_{11}$. We set:

$$Z(2, 3) = \frac{-(p_{31} - p_{11})^2 + 1 - p_{21} + p_{11} - (p_{23} - p_{33})^+}{p_{21} - p_{31}} \geq \frac{p_{31} - p_{11}}{p_{21} - p_{31}} Z(1, 3).$$

As in sub-case (B4), this choice is such that Claim 17 holds with equality. It just remains to verify the positivity of ν . As $Z(1, 3)$, $Z(2, 3)$ and $Z(3, 2)$ are the only possible negative entries in τ , we only need to check for $\nu(\cdots, 3, 1, 3)$, $\nu(\cdots, 3, 2, 3)$, $\nu(\cdots, 3, 3, 2)$. By construction, $\tau(s = 2, a = 3 | s^{-1}, a^{-1}) \geq \frac{p_{31} - p_{11}}{p_{21} - p_{31}} \cdot \tau(s = 1, a = 3 | s^{-1}, a^{-1})$. Thus $\nu(\cdots, 3, 1, 3)$ is not a problem. As for $\nu(\cdots, 3, 2, 3)$, we have:

$$\begin{aligned} & \nu^*(a^{-k}, \dots, a^{-1} = 3, s = 2, a = 3) \\ &= \frac{p_{31} - p_{11}}{p_{21} - p_{31}} \cdot \nu^*(a^{-k}, \dots, a^{-1} = 3, s = 1, a^{-1} = 3) + \Delta \cdot \nu^*(a^{-k}, \dots, a^{-2}, s^{-1} = 1, a^{-1} = 3) \geq 0. \end{aligned}$$

Note that we define $\Delta = \frac{1 - p_{21} + p_{11} - (p_{23} - p_{33})^+}{p_{21} - p_{31}}$. Finally, we deal with $\nu(\cdots, 3, 3, 2)$. From the recursive formula for ν , we obtain:

$$\begin{aligned} & \nu(a^{-k}, \dots, a^{-1} = 3, s = 3, a = 2) - \nu(a^{-k}, \dots, a^{-1} = 3, s = 2, a = 3) \\ &= (Z(3, 2) - Z(2, 3)) \cdot \nu(\cdots, s^{-1} = 1, a^{-1} = 3) + (U(3, 2) - U(2, 3)) \cdot \nu(\cdots, s^{-1} = 2, a^{-1} = 3) \\ &= (p_{32} - p_{22}) \cdot \nu(a^{-k}, \dots, a^{-2}, s^{-1} = 2, a^{-1} = 3) - (p_{12} - p_{32}) \cdot \nu(a^{-k}, \dots, a^{-2}, s^{-1} = 1, a^{-1} = 3). \end{aligned}$$

Combining the previous two equations, we deduce:

$$\begin{aligned} & \nu(a^{-k}, \dots, a^{-1} = 3, s = 3, a = 2) \\ & \geq (p_{32} - p_{22}) \cdot \nu(a^{-k}, \dots, a^{-2}, s^{-1} = 2, a^{-1} = 3) - (p_{12} - p_{32} - \Delta) \cdot \nu(a^{-k}, \dots, a^{-2}, s^{-1} = 1, a^{-1} = 3). \end{aligned}$$

We already know $\nu(\cdots, 2, 3) \geq \frac{p_{31} - p_{11}}{p_{21} - p_{31}} \cdot \nu(\cdots, 1, 3)$. Thus to show $\nu(a^{-k}, \dots, a^{-1} = 3, s = 3, a = 2) \geq 0$ it suffices to show:

$$p_{32} - p_{22} \geq \frac{p_{21} - p_{31}}{p_{31} - p_{11}} \cdot (p_{12} - p_{32} - \Delta).$$

Plugging in the definition of Δ , we need to show:

$$(p_{32} - p_{22})(p_{31} - p_{11}) + (1 - p_{21} + p_{11} - (p_{23} - p_{33})^+) \geq (p_{12} - p_{32})(p_{21} - p_{31}).$$

If $p_{23} < p_{33}$, then the above inequality holds because $1 - \beta \geq \beta^2$. Otherwise we need to show $(p_{32} - p_{22})(p_{31} - p_{11}) + p_{11} + p_{22} + p_{33} \geq (p_{12} - p_{32})(p_{21} - p_{31})$, which is equivalent to $(1 - p_{32} + p_{22})(1 - p_{31} + p_{11}) \geq (p_{12} - p_{32})(p_{21} - p_{31})$. This holds because $(1 - p_{32} + p_{22}) \geq p_{12} - p_{32}$

and $1 - p_{31} + p_{11} \geq p_{21} - p_{31}$.

- (d) $p_{22} > p_{32} > p_{12}$. We deduce from $p_{21} \geq p_{31}$ that $p_{23} < p_{33}$, and so $Y(2, 3) = 0$ from (71). Let us choose $Z(2, 3) = 0$. This way, τ only has one negative entry off the diagonal, and $\nu \geq 0$ holds by the usual argument. It remains to check Claim 17, which in this instance becomes:

$$\mu(1, 3) \cdot \max\left\{p_{22} - p_{12}, \frac{(p_{21} - p_{11})(p_{31} - p_{11})}{p_{21} - p_{31}}\right\} + \mu(2, 3) \cdot \max\{p_{22} - p_{32}, p_{31} - p_{11}\} \leq \mu(2, 3).$$

Because $\mu(1, 3) = (p_{21} - p_{31})\mu(3, 2)$ and $\mu(2, 3) = \mu(3, 2) - \mu(1, 3)$, the above reduces to:

$$\max\{(p_{22} - p_{12})(p_{21} - p_{31}), (p_{21} - p_{11})(p_{31} - p_{11})\} \leq (1 - p_{21} + p_{31}) \cdot \min\{1 - p_{22} + p_{32}, 1 - p_{31} + p_{11}\}. \quad (73)$$

Suppose the max and min operators both select the first term. Then (73) becomes $(p_{22} - p_{12})(p_{21} - p_{31}) \leq (1 - p_{21} + p_{31})(1 - p_{22} + p_{32})$, which follows from term-wise comparisons.

Suppose the max is achieved by the first term while the min is achieved by the second term, then we need to show $(p_{22} - p_{12})(p_{21} - p_{31}) \leq (1 - p_{21} + p_{31})(1 - p_{31} + p_{11})$. This is a quadratic inequality in p_{31} with negative leading coefficient, and it holds at the extreme values $p_{31} = p_{11}$ and $p_{31} = p_{21}$.

Next suppose the max in (73) is achieved by the second term while the min is achieved by the first. Then we need to show $(p_{21} - p_{11})(p_{31} - p_{11}) \leq (1 - p_{21} + p_{31})(1 - p_{22} + p_{32})$. This is linear in p_{31} and holds at both $p_{31} = p_{11}$ and $p_{31} = p_{21}$.

Finally, suppose the max and the min are achieved by the second term. We need to show $(p_{21} - p_{11})(p_{31} - p_{11}) \leq (1 - p_{21} + p_{31})(1 - p_{31} + p_{11})$. This is quadratic in p_{31} and holds at the extreme values.

This tricky sub-case is at last resolved, and with that we are done analyzing the case of two type-0 constraints plus a type-I constraint.

G.6 One type-0 constraint

In this subsection, we study those μ having only one zero off the diagonal. Thus there are two type-I constraints that bind at both μ and μ^{tt} . Up to permutation, we could assume that either they are both centered at state 3, or they are centered at states 2 and 3. When the

former happens, we know from Claim 8 that ν^* has configuration (A3)+(A4) (the situation with (A5)+(A6) being symmetric), and the binding type-0 constraint must be $\mu(3,1) = 0$ or $\mu(3,2) = 0$. When the latter happens, we know from Claim 9 that ν^* has configuration (A7), and the type-0 constraint is $\mu(2,3) = 0$ or $\mu(3,2) = 0$. These generate four possible configurations for ν^* :

$$\nu^* = \begin{pmatrix} + & + & + \\ + & + & + \\ 0 & 0 & + \end{pmatrix} \begin{pmatrix} + & + & 0 \\ + & + & 0 \\ 0 & + & + \end{pmatrix} \begin{pmatrix} + & + & 0 \\ + & + & 0 \\ 0 & 0 & + \end{pmatrix} \quad (\text{B9})$$

$$\nu^* = \begin{pmatrix} + & + & + \\ + & + & + \\ 0 & 0 & + \end{pmatrix} \begin{pmatrix} + & + & 0 \\ + & + & 0 \\ + & 0 & + \end{pmatrix} \begin{pmatrix} + & + & 0 \\ + & + & 0 \\ 0 & 0 & + \end{pmatrix} \quad (\text{B10})$$

$$\nu^* = \begin{pmatrix} + & 0 & + \\ + & + & 0 \\ 0 & 0 & + \end{pmatrix} \begin{pmatrix} + & + & 0 \\ 0 & + & 0 \\ + & + & + \end{pmatrix} \begin{pmatrix} + & + & 0 \\ 0 & + & 0 \\ + & + & + \end{pmatrix} \quad (\text{B11})$$

$$\nu^* = \begin{pmatrix} + & 0 & + \\ + & + & + \\ 0 & 0 & + \end{pmatrix} \begin{pmatrix} + & + & 0 \\ 0 & + & 0 \\ + & 0 & + \end{pmatrix} \begin{pmatrix} + & + & 0 \\ 0 & + & 0 \\ + & 0 & + \end{pmatrix} \quad (\text{B12})$$

G.6.1 Sub-case (B9)

The binding dual constraints imply the following about μ :

$$\mu(3,1) = 0.$$

$$\mu(3,2) = \mu(2,1)(p_{13} - p_{23}).$$

$$\mu(3,2)(1 - p_{33} + p_{23}) = \mu(1,2)(p_{13} - p_{23}).$$

$$\mu(3,2)(p_{33} - p_{23}) = \mu(1,3)(p_{13} - p_{23}).$$

$$\mu(3,2)(p_{13} - p_{33}) = \mu(2,3)(p_{13} - p_{23}).$$

Note that the fourth equation follows from the previous two and $\mu(2,1) = \mu(1,2) + \mu(1,3)$, while the last equation follows from the fourth and $\mu(3,2) = \mu(1,3) + \mu(2,3)$. From these we obtain $p_{13} > p_{33} > p_{23}$.

We need to determine $\tau(s, a|s^{-1}, a^{-1})$ for $(s^{-1}, a^{-1}) = (2, 1), (1, 2), (3, 2), (1, 3), (2, 3)$. Again

we follow the general principle to set $\tau(s, a|s^{-1}, a^{-1}) = 0$ whenever $\nu^*(a^{-1}, s, a) = 0$. However, there is a *caveat* here, because for $(s^{-1}, a^{-1}) = (1, 3)$ or $(2, 3)$, this principle is inconsistent with the desired row and column sums. We get around this issue by writing $\nu^*(3, 3, 2) = 0 = \mu(1, 3)(p_{13} - p_{33}) + \mu(2, 3)(p_{23} - p_{33})$, which suggests $\tau(3, 2|1, 3) = p_{13} - p_{33}$ and $\tau(3, 2|2, 3) = p_{23} - p_{33}$. The following summarizes the partially-determined τ :

?	?	$p_{33} - p_{23}$?	?	0	?	?	0	?	?	0	?	?	0
?	?	$p_{13} - p_{33}$?	?	0	?	?	0	?	?	0	?	?	0
0	0	p_{23}	0	$p_{13} - p_{23}$	p_{23}	0	$p_{33} - p_{23}$	p_{23}	0	$p_{13} - p_{33}$	p_{33}	0	$p_{23} - p_{33}$	p_{33}
$s^{-1}=2, a^{-1}=1$			$s^{-1}=1, a^{-1}=2$			$s^{-1}=3, a^{-1}=2$			$s^{-1}=1, a^{-1}=3$			$s^{-1}=2, a^{-1}=3$		

(74)

We have chosen $\tau(1, 3|2, 1) = p_{33} - p_{23}$ to be consistent with condition (41) and $\mu(1, 3) = (p_{33} - p_{23})\mu(2, 1)$. Furthermore, even though $\tau(3, 2|2, 3) = p_{23} - p_{33} < 0$, this does not affect the positivity of ν because for every s^{-1}, a^{-1} :

$$\tau(1, 3|s^{-1}, a^{-1})(p_{13} - p_{33}) + \tau(2, 3|s^{-1}, a^{-1})(p_{23} - p_{33}) = 0. \quad (75)$$

This implies $\nu(a^{-k}, \dots, a^{-1} = 3, s = 3, a = 2) = 0$ via the recursive formula (45), which is just what we expect from $\nu^*(3, 3, 2) = 0$.

Let us call the five matrices in (74) X, Y, Z, U, V . To have the correct row and column sums, we would like to have $X(1, 2) - X(2, 1) = p_{21} - p_{11} - (p_{33} - p_{23}) = 1 - p_{11} - p_{22} - p_{33}$. To ensure positivity, we set $X(1, 2) = (1 - p_{11} - p_{22} - p_{33})^+$. Similar analysis for Y and Z suggests the choices:

$$\begin{aligned} X(1, 2) &= (1 - p_{11} - p_{22} - p_{33})^+, \\ Y(1, 2) &= (p_{11} - p_{21})^+, \\ Z(1, 2) &= (p_{31} - p_{21})^+. \end{aligned} \quad (76)$$

For the matrices U and V , we normally would choose $U(1, 2) = (p_{11} - p_{31})^+$ and $V(1, 2) = (p_{21} - p_{31})^+$ so that ν is positive. However, this is more than necessary. Due to the exact proportionality (75), it suffices to have $(p_{33} - p_{23})U(1, 2) + (p_{13} - p_{33})V(1, 2)$ and $(p_{33} - p_{23})U(2, 1) + (p_{13} - p_{33})V(2, 1)$ both be positive. Thus we impose:

$$(p_{33} - p_{23})U(1, 2) + (p_{13} - p_{33})V(1, 2) = [(p_{33} - p_{23})(p_{11} - p_{31}) + (p_{13} - p_{33})(p_{21} - p_{31})]^+. \quad (77)$$

The τ constructed this way has the correct row and column sums, and it also ensures $\nu \geq 0$ by induction. Thus it remains to check the following analogue of Claim 17:

Claim 18 *Given $p_{ij} \leq \beta = \frac{\sqrt{5}-1}{2}$, it holds that:*

$$\mu(2, 1) \cdot X(1, 2) + \mu(1, 2) \cdot Y(1, 2) + \mu(3, 2) \cdot Z(1, 2) + \mu(1, 3) \cdot U(1, 2) + \mu(2, 3) \cdot V(1, 2) \leq \mu(1, 2).$$

Proof. We discuss three possibilities:

1. $p_{11}, p_{31} \geq p_{21}$ or $p_{11}, p_{31} \leq p_{21}$. Here we have constant signs at $a^{-1} = 1$ (trivially) and $a^{-1} = 2$. Constant signs is also guaranteed at $a^{-1} = 3$ because only the linear sum $(p_{33} - p_{23})U(\cdot, \cdot) + (p_{13} - p_{33})V(\cdot, \cdot)$ matters. The claim follows from $\mu \in \mathcal{M}_1$.
2. $p_{11} > p_{21} > p_{31}$. Plugging in (76) and (77) and expressing everything in terms of $\mu(2, 1)$, we can simplify the claim to:

$$(1 - p_{11} - p_{22} - p_{33})^+ + (p_{11} - p_{21}) + (p_{21} - p_{31})(p_{13} - p_{23}) \leq 1 - p_{33} + p_{23}.$$

This is equivalent to $(1 - p_{11} - p_{22} - p_{33})^+ + (p_{21} - p_{31})(p_{13} - p_{23}) \leq 2 - p_{11} - p_{22} - p_{33}$, which obviously holds when $p_{11} + p_{22} + p_{33} \leq 1$. If instead $p_{11} + p_{22} + p_{33} > 1$, we need to show $(p_{21} - p_{31})(p_{13} - p_{23}) \leq 2 - p_{11} - p_{22} - p_{33}$, which holds because the L.H.S. is at most $(1 - p_{22})(1 - p_{11})$ and $p_{11}p_{22} \leq 1 - p_{33}$.

3. $p_{31} > p_{21} > p_{11}$. Then $Y(1, 2) = 0$ and $\mu(1, 3) \cdot U(1, 2) + \mu(2, 3) \cdot V(1, 2) = 0$. The claim then reduces to:

$$(1 - p_{11} - p_{22} - p_{33})^+ + (p_{31} - p_{21})(p_{13} - p_{23}) \leq 1 - p_{33} + p_{23}.$$

When $p_{11} + p_{22} + p_{33} < 1$, the above is equivalent to $(p_{31} - p_{21})(p_{13} - p_{23}) \leq p_{11} + p_{22} + p_{23} = 1 - p_{21} + p_{11}$. Otherwise the above is equivalent to $(p_{31} - p_{21})(p_{13} - p_{23}) \leq 1 - p_{33} + p_{23}$. Both follow from $\beta^2 \leq 1 - \beta$.

We have thus proved the claim and verified $\mu \in \overline{\mathcal{C}}_*$ in this sub-case. ■

G.6.2 Sub-case (B10)

Here we can obtain the following conditions for μ :

$$\begin{aligned}\mu(3, 2) &= 0. \\ \mu(3, 1) &= \mu(1, 2)(p_{13} - p_{23}). \\ \mu(3, 1)(1 - p_{13} + p_{33}) &= \mu(2, 1)(p_{13} - p_{23}). \\ \mu(3, 1)(p_{13} - p_{33}) &= \mu(2, 3)(p_{13} - p_{23}). \\ \mu(3, 1)(p_{33} - p_{23}) &= \mu(1, 3)(p_{13} - p_{23}).\end{aligned}$$

As in the previous sub-case, we also have $p_{13} > p_{33} > p_{23}$. To determine τ , we run into a similar *caveat*, which we can resolve by writing $\nu^*(3, 3, 1) = \mu(1, 3)(p_{13} - p_{33}) + \mu(2, 3)(p_{23} - p_{33})$. The conditional τ can be partially determined as follows:

?	?	$p_{33} - p_{23}$?	?	$\frac{(p_{13}-p_{33})(p_{33}-p_{23})}{p_{13}-p_{23}}$?	?	0	?	?	0	?	?	0
?	?	$p_{13} - p_{33}$?	?	$\frac{(p_{13}-p_{33})^2}{p_{13}-p_{23}}$?	?	0	?	?	0	?	?	0
0	0	p_{23}	0	0	p_{33}	$p_{13} - p_{23}$	0	p_{23}	$p_{13} - p_{33}$	0	p_{33}	$p_{23} - p_{33}$	0	p_{33}
$s^{-1}=2, a^{-1}=1$			$s^{-1}=3, a^{-1}=1$			$s^{-1}=1, a^{-1}=2$			$s^{-1}=1, a^{-1}=3$			$s^{-1}=2, a^{-1}=3$		

(78)

A few comments are in order. First, as $\mu(3, 2) = 0$ in the current sub-case, the relevant (s^{-1}, a^{-1}) listed here are different from the previous sub-case. But we will still call these matrices X, Y, Z, U, V . Secondly, we have chosen $X(1, 3), X(2, 3), Y(1, 3), Y(2, 3)$ so as to maintain the exact proportionality:

$$\tau(1, 3|s^{-1}, a^{-1})(p_{13} - p_{33}) + \tau(2, 3|s^{-1}, a^{-1})(p_{23} - p_{33}) = 0. \quad (75)$$

These values also satisfy condition (41) for $(s, a) = (1, 3), (2, 3)$, as can be directly checked from the characterization of μ . It remains to fill out the question marks. We routinely choose:

$$\begin{aligned}X(1, 2) &= (1 - p_{11} - p_{22} - p_{33})^+. \\ Y(1, 2) &= \left(p_{31} - p_{11} - \frac{(p_{13} - p_{33})(p_{33} - p_{23})}{p_{13} - p_{23}} \right)^+. \\ Z(1, 2) &= (p_{22} - p_{12})^+.\end{aligned} \quad (79)$$

As in the previous sub-case, for U and V we simply require that:

$$(p_{33} - p_{23})U(1, 2) + (p_{13} - p_{33})V(1, 2) = [(p_{33} - p_{23})(p_{32} - p_{12}) + (p_{13} - p_{33})(p_{32} - p_{22})]^+. \quad (80)$$

Such a τ has the correct row and column sums, and the resulting ν is positive. It thus remains to verify the following analogue of Claim 18:

Claim 19 *Given $p_{ij} \leq \beta = \frac{\sqrt{5}-1}{2}$, it holds that:*

$$\mu(2, 1) \cdot X(1, 2) + \mu(3, 1) \cdot Y(1, 2) + \mu(1, 2) \cdot Z(1, 2) + \mu(1, 3) \cdot U(1, 2) + \mu(2, 3) \cdot V(1, 2) \leq \mu(1, 2).$$

Proof. There are again three possibilities:

1. $1 - p_{11} - p_{22} - p_{33}$ and $p_{31} - p_{11} - \frac{(p_{13}-p_{33})(p_{33}-p_{23})}{p_{13}-p_{23}}$ have the same sign. Then the claim follows from $\mu \in \mathcal{M}_1$.
2. $1 - p_{11} - p_{22} - p_{33} > 0 > p_{31} - p_{11} - \frac{(p_{13}-p_{33})(p_{33}-p_{23})}{p_{13}-p_{23}}$. Then the claim becomes:

$$\mu(2, 1)(1 - p_{11} - p_{22} - p_{33}) + \mu(1, 2)(p_{22} - p_{12})^+ + [\mu(1, 3)(p_{32} - p_{12}) + \mu(2, 3)(p_{32} - p_{22})]^+ \leq \mu(1, 2).$$

Since $\mu(2, 1) \leq \mu(1, 2)$, this obviously holds when the second “+” operator evaluates to zero. Below we assume this does not happen. If $p_{22} \geq p_{12}$, the above inequality simplifies to $(p_{13} - p_{33})(1 + p_{11} - p_{31}) + (p_{33} - p_{23})(p_{32} - p_{12}) \leq 1$ after expressing everything in terms of $\mu(1, 2)$. This last inequality is equivalent to $(p_{13} - p_{23})(p_{32} - p_{12}) \leq 1 - p_{13} + p_{33} + (p_{13} - p_{33})^2$, which follows from $\beta^2 \leq 1 - \beta$.

If instead $p_{22} < p_{12}$, the desired inequality becomes $(p_{13} - p_{23})(p_{32} - p_{12}) \leq 1 - p_{13} + p_{33} + (p_{13} - p_{33})^2 + p_{22} - p_{12} = p_{11} + p_{22} + p_{33} + (p_{13} - p_{33})^2$. This holds because $(p_{13} - p_{33})(p_{32} - p_{12}) = (p_{13} - p_{33})^2 + (p_{13} - p_{33})(p_{11} - p_{31}) \leq (p_{13} - p_{33})^2 + p_{11}$, and $(p_{33} - p_{23})(p_{32} - p_{12}) \leq p_{33}$.

3. $p_{31} - p_{11} - \frac{(p_{13}-p_{33})(p_{33}-p_{23})}{p_{13}-p_{23}} > 0 > 1 - p_{11} - p_{22} - p_{33}$. Here the claim becomes:

$$\begin{aligned} & \mu(3, 1) \cdot \left(p_{31} - p_{11} - \frac{(p_{13} - p_{33})(p_{33} - p_{23})}{p_{13} - p_{23}} \right) + \mu(1, 2)(p_{22} - p_{12})^+ \\ & + [\mu(1, 3)(p_{32} - p_{12}) + \mu(2, 3)(p_{32} - p_{22})]^+ \leq \mu(1, 2). \end{aligned}$$

If the second “+” operator evaluates to zero, this follows easily from $\mu(3, 1) = (p_{13} - p_{23})\mu(1, 2)$ and $\beta^2 \leq 1 - \beta$. Assume otherwise. If $p_{22} \geq p_{12}$, then we need to show $(1 - p_{13} + p_{33})(p_{22} - p_{12}) \leq 1 - (p_{13} - p_{33})^2$, which after factoring out $1 - p_{13} + p_{33}$ becomes $p_{22} - p_{12} \leq 1 + p_{13} - p_{33}$, or $p_{11} + p_{22} + p_{33} \leq 2$.

If instead $p_{22} < p_{12}$, then we need to show $(p_{13} - p_{33})(p_{12} - p_{22}) \leq 1 - (p_{13} - p_{33})^2$, which holds because $2\beta^2 \leq 1$.

The proof of the claim is completed. ■

G.6.3 Sub-case (B11)

From the proof of Claim 9, we have $\mu(2, 3) = 0$ and

$$\mu(1, 3) + \mu(2, 3) = \mu(2, 1)(p_{13} - p_{23}) + \mu(3, 1)(p_{13} - p_{33}). \quad (81)$$

$$\mu(1, 2) + \mu(3, 2) = \mu(3, 1)(p_{32} - p_{12}) + \mu(2, 1)(p_{22} - p_{12}). \quad (82)$$

Let us first use these equations to show:

$$p_{13} > p_{23}, p_{33}; \quad p_{22}, p_{32} > p_{12}.$$

Since $\mu(1, 3) = \mu(1, 3) + \mu(2, 3) = \mu(3, 1) + \mu(3, 2)$, we deduce $p_{13} > p_{23}$ from (81). Similarly from (82) we have $p_{32} > p_{12}$. $p_{13} > p_{33}$ follows by considering $\beta \cdot (81) + (82)$, and using $\mu(1, 3) > \mu(3, 1), \mu(1, 2) + \mu(3, 2) = \mu(2, 1)$. We can deduce $p_{22} > p_{12}$ likewise.

The configuration (B11) leads us to set the conditional τ partially as follows:

?	0	$p_{13} - p_{23}$?	0	$p_{13} - p_{33}$?	?	0	?	?	0			
$p_{22} - p_{12}$	p_{12}	0	$p_{32} - p_{12}$	p_{12}	0	0	p_{12}	0	0	p_{32}	0			
0	0	p_{23}	0	0	p_{33}	?	?	p_{23}	?	?	p_{23}			
$s^{-1}=2, a^{-1}=1$			$s^{-1}=3, a^{-1}=1$			$s^{-1}=1, a^{-1}=2$			$s^{-1}=3, a^{-1}=2$			$s^{-1}=1, a^{-1}=3$		

(83)

Next, with some foresight, we impose the following proportionality condition:

$$\tau(s = 3, a = 2 | s^{-1}, a^{-1}) = \frac{\mu(3, 2)}{\mu(1, 2)} \cdot \tau(s = 1, a = 2 | s^{-1}, a^{-1}). \quad (84)$$

Given this requirement and correct row and column sums, the last three matrices in (83) are

uniquely completed in the following way:

?	$\frac{\mu(1,2)}{\mu(2,1)}(p_{22} - p_{12})$	0
0	p_{12}	0
$p_{13} - p_{23} - \frac{\mu(3,2)}{\mu(2,1)}(p_{22} - p_{12})$	$\frac{\mu(3,2)}{\mu(2,1)}(p_{22} - p_{12})$	p_{23}

$$s^{-1}=1, a^{-1}=2$$

?	$\frac{\mu(1,2)}{\mu(2,1)}(p_{22} - p_{32})$	0
0	p_{32}	0
$p_{33} - p_{23} - \frac{\mu(3,2)}{\mu(2,1)}(p_{22} - p_{32})$	$\frac{\mu(3,2)}{\mu(2,1)}(p_{22} - p_{32})$	p_{23}

$$s^{-1}=3, a^{-1}=2$$

(85)

?	$\frac{\mu(1,2)}{\mu(2,1)}(p_{32} - p_{12})$	0
0	p_{12}	0
$p_{13} - p_{33} - \frac{\mu(3,2)}{\mu(2,1)}(p_{32} - p_{12})$	$\frac{\mu(3,2)}{\mu(2,1)}(p_{32} - p_{12})$	p_{33}

$$s^{-1}=1, a^{-1}=3$$

It remains to check $\nu \geq 0$.⁴⁴ As seen from (83), $\tau(s, a|s^{-1}, a^{-1})$ is positive when $a^{-1} = 1$ and $s \neq a$. Thus, by induction we have $\nu(\dots)$ positive off the diagonal for $a^{-1} = 1$.

We claim that the same is true for $a^{-1} = 3$. Specifically, we need to show $p_{13} - p_{33} \geq \frac{\mu(3,2)}{\mu(2,1)}(p_{32} - p_{12})$. From (82) and $\mu(2, 3) = 0$, we have $\mu(2, 1)(1 + p_{12} - p_{22}) = \mu(3, 1)(p_{32} - p_{12})$. From (81) we also have $\mu(3, 2) = \mu(1, 3) + \mu(2, 3) - \mu(3, 1) = \mu(2, 1)(p_{13} - p_{23}) - \mu(3, 1)(1 - p_{13} + p_{33})$. Thus we can compute that:

$$\begin{aligned} (p_{32} - p_{12}) \frac{\mu(3, 2)}{\mu(2, 1)} &= (p_{32} - p_{12})(p_{13} - p_{23}) - (1 + p_{12} - p_{22})(1 - p_{13} + p_{33}) \\ &= (p_{32} - p_{12})(p_{13} - p_{23}) - (p_{22} - p_{12})(p_{13} - p_{33}) - (1 + p_{12} - p_{22}) + (p_{13} - p_{33}) \\ &\leq p_{13} - p_{33}. \end{aligned}$$

The last inequality is because $\beta^2 \leq 1 - \beta$. This is exactly what we wanted to show, and so $\nu(\dots) \geq 0$ for $a^{-1} = 3$ as well.

Finally we need to prove $\nu(a^{-k}, \dots, a^{-1}, s, a) \geq 0$ for $a^{-1} = 2$ and $s \neq a$. Thanks to the proportionality condition (84), we simply need to verify:⁴⁵

$$\mu(1, 2) \cdot \tau(s, a|s^{-1} = 1, a^{-1} = 2) + \mu(3, 2) \cdot \tau(s, a|s^{-1} = 3, a^{-1} = 2) \geq 0, \forall (s, a) = (1, 2), (3, 2), (3, 1). \quad (86)$$

For $(s, a) = (1, 2), (3, 2)$, the above is equivalent to $\mu(1, 2)(p_{22} - p_{12}) + \mu(3, 2)(p_{22} - p_{32}) \geq 0$,

⁴⁴The proportionality condition guarantees $\sum_{s^{-1}, a^{-1}} \tau(s^{-1}, a^{-1}, s, a) = \mu(s, a)$ for $(s, a) = (1, 2), (3, 2)$. This then holds for every (s, a) by the other requirements on τ and linear dependence.

⁴⁵The proportionality condition implies that for all a^{-k}, \dots, a^{-2} , the ratio between $\nu(a^{-k}, \dots, a^{-2}, s^{-1} = 1, a^{-1} = 2)$ and $\nu(a^{-k}, \dots, a^{-1}, s^{-1} = 1, a^{-1} = 2)$ is equal to the ratio between $\mu(1, 2)$ and $\mu(3, 2)$.

which is equivalent to $\mu(3, 2)(p_{32} - p_{12})\mu(2, 1)(p_{22} - p_{12}) \geq \mu(3, 2)(p_{32} - p_{12})$. Similar to the computation above, we do have:

$$\begin{aligned} (p_{32} - p_{12})\frac{\mu(3, 2)}{\mu(2, 1)} &= (p_{13} - p_{23})(p_{32} - p_{12}) - (1 + p_{12} - p_{22})(1 - p_{13} + p_{33}) \\ &= (p_{13} - p_{23})(p_{32} - p_{12}) - (p_{22} - p_{12})(p_{13} - p_{33}) - (1 - p_{13} + p_{33}) + (p_{22} - p_{12}) \\ &\leq p_{22} - p_{12}. \end{aligned}$$

To prove (86) for $(s, a) = (3, 1)$, let us denote the three matrices in (85) by X , Y and Z . We have $\mu(1, 2)X(3, 1) + \mu(3, 2)Y(3, 1) + \mu(1, 3)Z(3, 1) = \mu(3, 1)$. It thus suffices to show $\mu(1, 3)Z(3, 1) \leq \mu(3, 1)$. From (81) and (82), we have $\frac{\mu(2, 1)}{\mu(3, 1)} = \frac{p_{32} - p_{12}}{1 + p_{12} - p_{22}}$ and:

$$\frac{\mu(1, 3)}{\mu(3, 1)} = \frac{(p_{32} - p_{12})(p_{13} - p_{23})}{1 + p_{12} - p_{22}} + p_{13} - p_{33} \leq \frac{\beta^2}{1 - \beta} + \beta = 1 + \beta.$$

Note that $Z(3, 1) \leq p_{13} - p_{33} \leq \beta$, and so $\frac{\mu(1, 3)}{\mu(3, 1)}Z(3, 1) \leq \beta(1 + \beta) = 1$ as desired.

Thus the τ we construct satisfy all the conditions for $\mu \in \overline{\mathcal{C}}_*$.

G.6.4 Sub-case (B12)

Similar to the previous sub-case, we can again show $p_{13} > p_{23}, p_{33}$ and $p_{22}, p_{32} > p_{12}$. The conditional τ is partially determined as follows:

?	0	?	?	0	?	?	$p_{22} - p_{12}$	0	?	$p_{32} - p_{12}$	0	?	$p_{32} - p_{22}$	0
?	p_{12}	?	?	p_{12}	?	0	p_{12}	0	0	p_{12}	0	0	p_{22}	0
0	0	p_{23}	0	0	p_{33}	$p_{13} - p_{23}$	0	p_{23}	$p_{13} - p_{33}$	0	p_{33}	$p_{23} - p_{33}$	0	p_{33}
$s^{-1}=2, a^{-1}=1$			$s^{-1}=3, a^{-1}=1$			$s^{-1}=1, a^{-1}=2$			$s^{-1}=1, a^{-1}=3$			$s^{-1}=2, a^{-1}=3$		

(87)

So far, we have the correct row and column sums. Moreover, τ is positive off the diagonal when $a^{-1} = 2$. Thus $\nu \geq 0$ for $a^{-1} = 2$. However, for $a^{-1} = 3$, we see that $\tau(s = 1, a = 2 | s^{-1} = 2, a^{-1} = 3) = p_{32} - p_{22}$ and $\tau(s = 2, a = 3 | s^{-1} = 2, a^{-1} = 3) = p_{23} - p_{33}$ are possibly negative. To ensure the positivity of $\nu(a^{-k}, \dots, a^{-1} = 3, s = 1, a = 2)$ and $\nu(a^{-k}, \dots, a^{-1} = 3, s = 3, a = 1)$, we impose the following proportionality condition:

$$\tau(s = 2, a = 3 | s^{-1}, a^{-1}) = \frac{\mu(2, 3)}{\mu(1, 3)} \cdot \tau(s = 1, a = 3 | s^{-1}, a^{-1}). \quad (88)$$

Under this condition, $\nu(a^{-k}, \dots, a^{-1} = 3, s = 1, a = 2) \geq 0$ if $\mu(1, 3)(p_{32} - p_{12}) + \mu(2, 3)(p_{32} - p_{22}) \geq 0$. Given configuration (B12), we do have $\mu(1, 3)(p_{32} - p_{12}) + \mu(2, 3)(p_{32} - p_{22}) = \nu^*(3, 1, 2) \geq 0$. This means $\nu \geq 0$ for $a^{-1} = 3$ as well.

The proportionality condition (88), together with row and column sums, implies that the first two matrices in (87) should be completed as follows:

?	0	$\frac{\mu(1,3)}{\mu(3,1)}(p_{13} - p_{23})$?	0	$\frac{\mu(1,3)}{\mu(3,1)}(p_{13} - p_{33})$
$p_{22} - p_{12} - \frac{\mu(2,3)}{\mu(3,1)}(p_{13} - p_{23})$	p_{12}	$\frac{\mu(2,3)}{\mu(3,1)}(p_{13} - p_{23})$	$p_{32} - p_{12} - \frac{\mu(2,3)}{\mu(3,1)}(p_{13} - p_{33})$	p_{12}	$\frac{\mu(2,3)}{\mu(3,1)}(p_{13} - p_{33})$
0	0	p_{23}	0	0	p_{33}
$s^{-1}=2, a^{-1}=1$			$s^{-1}=3, a^{-1}=1$		

(89)

It only remains to check the positivity of ν for $a^{-1} = 1$. We claim that in fact τ is positive off the diagonal when $a^{-1} = 1$:

$$\frac{\mu(2, 3)}{\mu(3, 1)}(p_{13} - p_{23}) \leq p_{22} - p_{12}.$$

$$\frac{\mu(2, 3)}{\mu(3, 1)}(p_{13} - p_{33}) \leq p_{32} - p_{12}.$$

Note that we still have the same two equations (81) and (82) from the previous sub-case, although now $\mu(3, 2) = 0$. From (82), $\mu(2, 3) = \mu(1, 2) + \mu(3, 2) - \mu(2, 1) = \mu(3, 1)(p_{32} - p_{12}) + \mu(2, 1)(p_{22} - p_{12} - 1) \leq \mu(3, 1)(p_{32} - p_{12})$. This implies the second inequality above.

From (81), we have $\mu(3, 1)(1 - p_{13} + p_{33}) = \mu(2, 1)(p_{13} - p_{23})$. Thus:

$$\begin{aligned} \mu(2, 3)(p_{13} - p_{23}) &= \mu(3, 1)(p_{13} - p_{23})(p_{32} - p_{12}) + \mu(2, 1)(p_{13} - p_{23})(p_{22} - p_{12} - 1) \\ &= \mu(3, 1) ((p_{13} - p_{23})(p_{32} - p_{12}) - (1 - p_{13} + p_{33})(1 - p_{22} + p_{12})). \end{aligned}$$

The first inequality above is then reduced to $(p_{13} - p_{23})(p_{32} - p_{12}) - (1 - p_{13} + p_{33})(1 - p_{22} + p_{12}) \leq p_{22} - p_{12}$. This is equivalent to $(p_{13} - p_{23})(p_{32} - p_{12}) \leq 1 - p_{13} + p_{33} + (p_{13} - p_{33})(p_{22} - p_{12})$, which holds because $\beta^2 \leq 1 - \beta$.

With that we have shown that the μ in this last sub-case also belongs to $\overline{\mathcal{C}}_*$. Hence we have also completed the long proof of Proposition 5' and of Proposition ??.

H Insufficiency of \mathcal{C}_k

In this section we prove that even with 3 states, no finite memory is generally sufficient for implementation:

Proposition 7' *For any k there exists $\varepsilon > 0$, such that if p is not pseudo-renewal, p has full-support and $p_{ii} > 1 - \varepsilon, \forall i$, then $C_1 \not\supseteq C_2 \not\supseteq \cdots \not\supseteq C_k \not\supseteq C_{k+1}$.*

This will immediately imply Proposition 6 stated in the main text. In what follows, we first provide a proof for $k = 1$ to illustrate the methods and then proceed to larger k .

H.1 Proof for $k = 1$

Since the process is not renewal, we may without loss assume:

$$p_{31} > p_{21}. \tag{90}$$

This assumption will be maintained throughout this appendix.

We are going to consider those μ that have the correct marginals and $\mu(1, 2) = \mu(3, 1) = 0$. These μ form a polygon on a 2-dimensional plane, thus the restriction of \mathcal{C}_k onto this plane is determined by two rays, which correspond to minimum and maximum values of the ratio $\frac{\mu(1,3)}{\mu(3,2)}$. To be a little more specific, all μ on this plane can be parameterized by:

$$\mu(s, a) = \begin{pmatrix} \lambda_1 - y & 0 & y \\ y & \lambda_2 - x & x - y \\ 0 & x & \lambda_3 - x \end{pmatrix}$$

with $x \geq y \geq 0$. It is then easy to see that a ray on this plane is determined by the ratio $\frac{\mu(1,3)}{\mu(3,2)}$.

Suppose $\nu(a^{-1}, s, a)$ is a solution to the linear system for $\mu \in \mathcal{M}_1$, then we have:

$$\mu(1, 3) \geq \nu(2, 1, 3). \tag{91}$$

Since $\nu(2, 1, 2) = \nu(2, 3, 1) = 0$, we compute that:

$$\begin{aligned}
\nu(2, 1, 3) &= [\nu(2, 1, 1) + \nu(2, 1, 2) + \nu(2, 1, 3)] - [\nu(2, 1, 1) + \nu(2, 2, 1) + \nu(2, 3, 1)] + \nu(2, 2, 1) \\
&= [\mu(2, 2)p_{21} + \mu(3, 2)p_{31}] - [(\mu(2, 2) + \mu(3, 2))p_{21}] + \nu(2, 2, 1) \\
&= \mu(3, 2)(p_{31} - p_{21}) + \nu(2, 2, 1).
\end{aligned} \tag{92}$$

Thus we deduce from (91) and (92) that a necessary condition for $\mu \in \mathcal{C}_1$ is

$$\mu(1, 3) \geq \mu(3, 2) \cdot (p_{31} - p_{21}), \forall \mu \in \mathcal{C}_1. \tag{93}$$

We will prove the following claim that implies $\mathcal{C}_1 \neq \mathcal{C}_2$:

Claim 20 *Suppose $p_{11} + p_{22} + p_{33} > 2$. Then there exists $\mu \in \mathcal{C}_1$ such that $\mu(1, 3) = \mu(3, 2)(p_{31} - p_{21})$. But for every $\mu \in \mathcal{C}_2$ it holds that $\mu(1, 3) > \mu(3, 2)(p_{31} - p_{21})$.*

Proof. Let us first handle the second part. Suppose $\nu(a^{-2}, a^{-1}, s, a)$ solves the linear system for \mathcal{M}_2 together with some μ . Define $\nu(a^{-1}, s, a)$ to be the marginal of ν in those coordinates, then we have:

$$\mu(1, 3) \geq \nu(2, 1, 3) + \nu(2, 3, 1, 3). \tag{94}$$

From (92) we have $\nu(2, 1, 3) \geq \mu(3, 2)(p_{31} - p_{21})$. On the other hand, from $\nu(2, 3, 1, 2) = \nu(2, 3, 3, 1) = 0$ we obtain:

$$\begin{aligned}
\nu(2, 3, 1, 3) &= \sum_a \nu(2, 3, 1, a) - \nu(2, 3, 1, 1) \\
&= \sum_a \nu(2, 3, 1, a) - \sum_s \nu(2, 3, s, 1) + \nu(2, 3, 2, 1) \\
&= \nu(2, 1, 3)(p_{11} - p_{31}) - \nu(2, 2, 3)(p_{31} - p_{21}) + \nu(2, 3, 2, 1).
\end{aligned} \tag{95}$$

Next we compute $\nu(2, 2, 3)$ as follows:

$$\begin{aligned}
\nu(2, 2, 3) &= \sum_a \nu(2, 2, a) - \nu(2, 2, 2) - \nu(2, 2, 1) \\
&= \sum_a \nu(2, 2, a) - \left(\sum_s \nu(2, s, 2) - \nu(2, 3, 2) \right) - \nu(2, 2, 1) \\
&= \mu(3, 2)(p_{32} - p_{22}) + \nu(2, 3, 2) - \nu(2, 2, 1) \\
&= \mu(3, 2)(1 - p_{22} + p_{32}) - \nu(1, 3, 2) - \nu(3, 3, 2) - \nu(2, 2, 1) \\
&\leq \mu(3, 2)(1 - p_{22} + p_{32}) - \nu(3, 3, 2).
\end{aligned} \tag{96}$$

Plugging (92) and (96) into (95), we obtain:

$$\nu(2, 3, 1, 3) \geq \mu(3, 2)(p_{31} - p_{21})(p_{11} - p_{31} - 1 + p_{22} - p_{32}) = \mu(3, 2)(p_{31} - p_{21}) \cdot (p_{11} + p_{22} + p_{33} - 2).$$

Inserting this into (94), we conclude that a necessary condition for $\mu \in \mathcal{C}_2$ is:

$$\mu(1, 3) \geq \mu(3, 2)(p_{31} - p_{21})(p_{11} + p_{22} + p_{33} - 1), \forall \mu \in \mathcal{C}_2. \tag{97}$$

By assumption $p_{11} + p_{22} + p_{33} - 1 > 1$, thus the above implies $\mu(1, 3) > \mu(3, 2)(p_{31} - p_{21})$, as we needed to show.⁴⁶

Next we show that $\mu(1, 3) = \mu(3, 2)(p_{31} - p_{21})$ is attainable for some $\mu \in \mathcal{C}_1$. For this we distinguish between two possibilities:

1. $p_{13} \geq p_{23}$. In order for inequality (93) to hold, we need $\nu(1, 1, 3) = \nu(3, 1, 3) = \nu(2, 2, 1) = 0$ from (91) and (92). We will solve for ν under the additional assumption that $\nu(1, 3, 2) = \nu(3, 3, 2) = 0$, then ν has the following configuration:

$$\nu = \begin{pmatrix} + & 0 & 0 \\ + & + & + \\ 0 & 0 & + \end{pmatrix} \begin{pmatrix} + & 0 & + \\ 0 & + & + \\ 0 & + & + \end{pmatrix} \begin{pmatrix} + & 0 & 0 \\ + & + & + \\ 0 & 0 & + \end{pmatrix}$$

⁴⁶The constant “2” appearing in the assumption $p_{11} + p_{22} + p_{33} > 2$ cannot be reduced. We can show (available from the authors) that if $p_{11} = p_{22} = p_{32} = 1 - 2w$ and other transition probabilities are equal to w , then $\mathcal{C}_1 = \mathcal{C}_k$ for every k . For such a non-renewal process, $p_{11} + p_{22} + p_{33} = 2 - 3w$ which can be made arbitrarily close to 2.

It follows that:

$$\nu(1, 1, 1) = \sum_a \nu(1, 1, a) = \mu(1, 1)p_{11} + \mu(2, 1)p_{21}.$$

$$\nu(1, 2, 1) = \sum_s \nu(1, s, 1) - \nu(1, 1, 1) = \mu(2, 1)(p_{11} - p_{21}).$$

$$\nu(1, 2, 2) = \sum_s \nu(1, s, 2) = \lambda_1 p_{12}.$$

$$\nu(1, 3, 3) = \sum_a \nu(1, 3, a) = \mu(1, 1)p_{13} + \mu(2, 1)p_{23}.$$

$$\nu(1, 2, 3) = \sum_s \nu(1, s, 3) - \nu(1, 3, 3) = \mu(2, 1)(p_{13} - p_{23}).$$

$$\nu(2, 1, 1) = \sum_s \nu(2, s, 1) = \lambda_2 p_{21}.$$

$$\nu(2, 1, 3) = \sum_{a^{-1}} \nu(a^{-1}, 1, 3) = \mu(1, 3).$$

$$\nu(2, 3, 2) = \sum_{a^{-1}} \nu(a^{-1}, 3, 2) = \mu(3, 2).$$

$$\nu(2, 2, 2) = \sum_s \nu(2, s, 2) - \nu(2, 3, 2) = \mu(2, 2)p_{22} - \mu(3, 2)(1 - p_{22}).$$

$$\nu(2, 2, 3) = \sum_a \nu(2, 2, a) - \nu(2, 2, 2) = \mu(2, 2)p_{22} + \mu(3, 2)p_{32} - \nu(2, 2, 2) = \mu(3, 2)(1 - p_{22} + p_{32}).$$

$$\nu(2, 3, 3) = \sum_a \nu(2, 3, a) - \nu(2, 3, 2) = \mu(2, 2)p_{23} - \mu(3, 2)(1 - p_{33}).$$

$$\nu(3, 1, 1) = \sum_a \nu(3, 1, a) = \mu(1, 3)p_{11} + \mu(2, 3)p_{21} + \mu(3, 3)p_{31}.$$

$$\nu(3, 2, 1) = \sum_{a^{-1}} \nu(a^{-1}, 2, 1) - \nu(1, 2, 1) = \mu(2, 1)(1 - p_{11} + p_{21}).$$

$$\nu(3, 2, 2) = \sum_s \nu(3, s, 2) = \lambda_3 p_{32}.$$

$$\nu(3, 3, 3) = \sum_a \nu(3, 3, a) = \mu(1, 3)p_{13} + \mu(2, 3)p_{23} + \mu(3, 3)p_{33}.$$

$$\nu(3, 2, 3) = \sum_s \nu(3, s, 3) - \nu(3, 3, 3) = \mu(1, 3)(p_{33} - p_{13}) + \mu(2, 3)(p_{33} - p_{23}).$$

Here $\nu(1, 2, 1) \geq 0$ because $p_{11} > 2 - p_{22} - p_{33} > 1 - p_{22} > p_{21}$. Also, $\nu(3, 2, 3) \geq 0$ because similarly $p_{33} \geq p_{13}, p_{23}$. While $\nu(2, 2, 2)$ and $\nu(2, 3, 3)$ could be negative, they do not cause any issue because they are on the diagonal; we can mix ν with ν^{tt} to make these entries positive. Thus ν is positive and $\mu \in \mathcal{C}_1$.

2. $p_{23} > p_{13}$. We still impose $\nu(1, 1, 3) = \nu(3, 1, 3) = \nu(2, 2, 1) = 0$. But here we can no longer support $\nu(1, 3, 2) = 0$, which would imply $\nu(1, 2, 3) = \mu(2, 1)(p_{13} - p_{23}) < 0$. Hence we assume $\nu(1, 2, 3) = \nu(3, 3, 2) = 0$ instead, leading to the following configuration of zeros:

$$\nu = \begin{pmatrix} + & 0 & 0 \\ + & + & 0 \\ 0 & + & + \end{pmatrix} \begin{pmatrix} + & 0 & + \\ 0 & + & + \\ 0 & + & + \end{pmatrix} \begin{pmatrix} + & 0 & 0 \\ + & + & + \\ 0 & 0 & + \end{pmatrix}$$

Again, we can solve ν in terms of μ . In the end, the results here differ from the above in that $\nu(1, 2, 3), \nu(1, 3, 2), \nu(2, 2, 2), \nu(3, 3, 3)$ are increased by $\Delta = \mu(2, 1)(p_{23} - p_{13})$, whereas $\nu(1, 2, 2), \nu(1, 3, 3), \nu(2, 2, 3), \nu(2, 3, 2)$ are decreased by Δ . We claim that the resulting ν is still positive. It suffices to check the latter four entries, which further reduces to checking the off-diagonal entries $\nu(2, 2, 3), \nu(2, 3, 2)$:

$$\nu(2, 2, 3) = \mu(3, 2)(1 - p_{22} + p_{32}) - \mu(2, 1)(p_{23} - p_{13}) = \mu(3, 2)(1 - p_{22} + p_{32}) - \mu(1, 3)(p_{23} - p_{13}).$$

$$\nu(2, 3, 2) = \mu(3, 2) - \mu(1, 3)(p_{23} - p_{13}).$$

Both are positive because $\mu(3, 2) = \mu(1, 3) + \mu(2, 3) \geq \mu(1, 3)$. Thus again $\mu \in \mathcal{C}_1$.

This completes the proof of the claim and of Proposition 7' for $k = 1$. ■

H.2 Interlude: tightness of $\frac{\sqrt{5}-1}{2}$

We have mentioned that the bound $\beta = \frac{\sqrt{5}-1}{2}$ in Proposition ?? cannot be improved upon. In this subsection we formally justify this:⁴⁷

Claim 21 *For any $\beta' > \beta$, there exists an open set of processes p with $0 < p_{ij} \leq \beta'$ and $\mathcal{C}_1 \neq \mathcal{C}_2$.*

⁴⁷A slightly less rigorous treatment was also provided as the last example in Section 3.4.2.

Proof. If $\beta' > \frac{2}{3}$, this follows easily from the last claim. To cover $\beta' > \beta$, we will refine the computation in the previous subsection. The key observation is that the last inequality in (96) holds equal only if $\nu(1, 3, 2) = 0$. But as we discussed, this is not possible when $p_{23} > p_{13}$. In fact we have:

$$\begin{aligned}
\nu(1, 3, 2) &= \sum_a \nu(1, 3, a) - \nu(1, 3, 3) \\
&= \mu(1, 1)p_{13} + \mu(2, 1) \cdot p_{23} - \sum_s \nu(1, s, 3) + \nu(1, 1, 3) + \nu(1, 2, 3) \\
&= \mu(2, 1)(p_{23} - p_{13}) + \nu(1, 1, 3) + \nu(1, 2, 3) \\
&= \mu(1, 3)(p_{23} - p_{13}) + \nu(1, 1, 3) + \nu(1, 2, 3).
\end{aligned}$$

Thus instead of (96), we obtain the stronger bound:

$$\nu(2, 2, 3) \leq \mu(3, 2)(1 - p_{22} + p_{32}) - \mu(1, 3)(p_{23} - p_{13}) - \nu(3, 3, 2). \quad (98)$$

Plugging (92) and (98) into (95) and then (94), we conclude the following strengthening of (97):

$$\mu(1, 3) \geq (p_{31} - p_{21})(\mu(3, 2)(p_{11} + p_{22} + p_{33} - 1) + \mu(1, 3)(p_{23} - p_{13})). \quad (99)$$

Simplifying, we obtain that a necessary condition for $\mu \in \mathcal{C}_2$ is:⁴⁸

$$\mu(1, 3) \geq \mu(3, 2)(p_{31} - p_{21}) \frac{p_{11} + p_{22} + p_{33} - 1}{1 - (p_{31} - p_{21})(p_{23} - p_{13})}.$$

If $p_{11} + p_{22} + p_{33} + (p_{31} - p_{21})(p_{23} - p_{13}) > 2$, then the above is more restrictive than $\mu(1, 3) \geq \mu(3, 2)(p_{31} - p_{21})$ and we will be able to deduce $\mathcal{C}_1 \neq \mathcal{C}_2$. It remains to find parameters $0 < p_{ij} \leq \beta'$ satisfying this inequality. Let $p_{11} = p_{22} = p_{33} = \beta'$, $p_{13} = p_{21} = p_{32} \rightarrow 0$ and $p_{12} = p_{23} = p_{31} \rightarrow 1 - \beta'$. Then indeed we have:

$$p_{11} + p_{22} + p_{33} + (p_{31} - p_{21})(p_{23} - p_{13}) \rightarrow 3\beta' + (1 - \beta')^2 = \beta'^2 + \beta' + 1 > 2.$$

This proves that the bound $\frac{\sqrt{5}-1}{2}$ on transition probabilities is tight in order for $\mathcal{C}_1 = \mathcal{C}_2$. ■

⁴⁸We will show in the next subsection that when $p_{13} \geq p_{23}$, there exists $\mu \in \mathcal{C}_2$ satisfying (97) with equality. When $p_{23} < p_{13}$, there exists $\mu \in \mathcal{C}_2$ satisfying (99) with equality. Thus the condition (97) or (99) characterize the minimum ratio $\frac{\mu(1,3)}{\mu(3,2)}$ for $\mu \in \mathcal{C}_2$, depending on whether p_{13} or p_{23} is larger. This dichotomy illustrates the subtlety in characterizing the set \mathcal{C}_k .

H.3 Proof for general k : lower bound on $\frac{\mu(1,3)}{\mu(3,2)}$

Here we develop the analogue of (97) or (99) for $\mu \in \mathcal{C}_k$. Suppose $\nu(a^{-k}, \dots, a^{-1}, s, a)$ together with μ solves the linear system defining \mathcal{M}_k . For $1 \leq m < k$, let $\nu(a^{-m}, \dots, a^{-1}, s, a)$ denote the marginal of ν in those coordinates. Note that $\nu(a^{-m}, \dots, a^{-1}, s, a)$ solves the linear system for \mathcal{M}_m .

By the positivity of ν , we must have:

$$\mu(1, 3) \geq \sum_{m=1}^k \nu(a^{-m} = 2, a^{-m-1} = \dots = a^{-1} = 3, s = 1, a = 3). \quad (100)$$

As before, we will give a lower bound to the R.H.S. sum in terms of μ . Starting from $\nu(2, 1, 3)$ and $\nu(2, 2, 3)$, we have the following recursive relations regarding $\nu(2, 3, \dots, 3, 1, 3)$ and $\nu(2, 3, \dots, 3, 2, 3)$:

$$\begin{aligned} & \nu(a^{-m} = 2, 3, \dots, 3, s = 1, a = 3) - \nu(a^{-m} = 2, 3, \dots, 3, s = 2, a = 1) \\ &= \sum_a \nu(a^{-m} = 2, 3, \dots, 3, s = 1, a) - \sum_s \nu(a^{-m} = 2, 3, \dots, 3, s, a = 1) \\ &= \nu(a^{-(m-1)} = 2, 3, \dots, 3, s = 1, a = 3)(p_{11} - p_{31}) - \nu(a^{-(m-1)} = 2, 3, \dots, 3, s = 2, a = 3)(p_{31} - p_{21}). \end{aligned} \quad (101)$$

$$\begin{aligned} & \nu(a^{-m} = 2, 3, \dots, 3, s = 2, a = 3) + \nu(a^{-m} = 2, 3, \dots, 3, s = 2, a = 1) - \nu(a^{-m} = 2, 3, \dots, 3, s = 3, a = 2) \\ &= \sum_a \nu(a^{-m} = 2, 3, \dots, 3, s = 2, a) - \sum_s \nu(a^{-m} = 2, 3, \dots, 3, s, a = 2) \\ &= \nu(a^{-(m-1)} = 2, 3, \dots, 3, s = 1, a = 3)(p_{12} - p_{32}) + \nu(a^{-(m-1)} = 2, 3, \dots, 3, s = 2, a = 3)(p_{22} - p_{32}). \end{aligned} \quad (102)$$

Simple induction leads to the following lemma:

Lemma 22 Define the first-order recursive sequences $\{\lambda_m\}$ and $\{\eta_m\}$:

$$\begin{aligned}
\lambda_1 &= 1; \\
\eta_1 &= 0; \\
\lambda_{m+1} &= (p_{11} - p_{31})\lambda_m - (p_{31} - p_{21})(p_{12} - p_{32})\eta_m; \\
\eta_{m+1} &= \lambda_m + (p_{22} - p_{32})\eta_m.
\end{aligned} \tag{103}$$

Suppose $\nu(a^{-k}, \dots, a^{-1}, s, a) \geq 0$ solves the linear system for \mathcal{M}_k , with $\mu(1, 2) = \mu(3, 1) = 0$. Then for $1 \leq m \leq k$ it holds that:

$$\begin{aligned}
&\nu(a^{-m} = 2, 3, \dots, 3, s = 1, a = 3) \\
&= \nu(2, 1, 3)\lambda_m - \nu(2, 2, 3)(p_{31} - p_{21})\eta_m \\
&\quad + \sum_{j=2}^m \nu(a^{-j} = 2, 3, \dots, 3, s = 2, a = 1)(\lambda_{m+1-j} + (p_{31} - p_{21})\eta_{m+1-j}) \\
&\quad - \sum_{j=2}^m \nu(a^{-j} = 2, 3, \dots, 3, s = 3, a = 2)(p_{31} - p_{21})\eta_{m+1-j}.
\end{aligned} \tag{104}$$

And similarly:

$$\begin{aligned}
&\nu(a^{-m} = 2, 3, \dots, 3, s = 2, a = 3) \\
&= \nu(2, 1, 3)(p_{12} - p_{32})\eta_m + \nu(2, 2, 3)(\lambda_m + (p_{22} - p_{32} - p_{11} + p_{31})\eta_m) \\
&\quad - \sum_{j=2}^m \nu(a^{-j} = 2, 3, \dots, 3, s = 2, a = 1)(\lambda_{m+1-j} + (p_{22} - p_{12} - p_{11} + p_{31})\eta_{m+1-j}) \\
&\quad + \sum_{j=2}^m \nu(a^{-j} = 2, 3, \dots, 3, s = 3, a = 2)(\lambda_{m+1-j} + (p_{22} - p_{32} - p_{11} + p_{31})\eta_{m+1-j}).
\end{aligned} \tag{105}$$

As a consequence of plugging (104) and (105) into the inequality (100), we obtain:

$$\begin{aligned}
\mu(1, 3) \geq & \nu(2, 1, 3) \cdot \sum_{j=1}^k \lambda_j - \nu(2, 2, 3) \cdot (p_{31} - p_{21}) \sum_{j=1}^k \eta_j \\
& + \sum_{j=2}^k \nu(a^{-j} = 2, 3, \dots, 3, s = 2, a = 1) \cdot \sum_{m=j}^k (\lambda_{m+1-j} + (p_{31} - p_{21}) \eta_{m+1-j}) \\
& - \sum_{j=2}^k \nu(a^{-j} = 2, 3, \dots, 3, s = 3, a = 2) \cdot (p_{31} - p_{21}) \sum_{m=j}^k \eta_{m+1-j}.
\end{aligned} \tag{106}$$

This enables us to deduce the following lower bounds on $\frac{\mu(1,3)}{\mu(3,2)}$:

Lemma 23 *Suppose $\lambda_j, \eta_j \geq 0, \forall 1 \leq j \leq k$. If $p_{13} \geq p_{23}$, then a necessary condition for $\mu \in \mathcal{C}_k$ and $\mu(1, 2) = \mu(3, 1) = 0$ is:*

$$\mu(1, 3) \geq \mu(3, 2) \cdot (p_{31} - p_{21}) \left[\sum_{j=1}^k \lambda_j - (1 - p_{22} + p_{32}) \sum_{j=1}^k \eta_j \right], \forall \mu \in \mathcal{C}_k. \tag{107}$$

If instead $p_{23} > p_{13}$, then the corresponding necessary condition is:

$$\mu(1, 3) \cdot \left(1 - (p_{31} - p_{21})(p_{23} - p_{13}) \sum_{j=1}^k \eta_j \right) \geq \mu(3, 2) \cdot (p_{31} - p_{21}) \left[\sum_{j=1}^k \lambda_j - (1 - p_{22} + p_{32}) \sum_{j=1}^k \eta_j \right], \forall \mu \in \mathcal{C}_k. \tag{108}$$

Proof. From (92) and (96) above, we have $\nu(2, 1, 3) \geq \mu(3, 2) \cdot (p_{31} - p_{21})$ and $\nu(2, 2, 3) \leq \mu(3, 2)(1 - p_{22} + p_{32}) - \nu(3, 3, 2) \leq \mu(3, 2) \cdot (1 - p_{22} + p_{32}) - \sum_{j=2}^k \nu(a^{-j} = 2, 3, \dots, 3, s = 3, a = 2)$. Plugging these into (106), we conclude (107) with extra positive terms on the R.H.S.

If $p_{23} > p_{13}$, we can use the stronger inequality (98) developed in the previous subsection. This enables us to conclude (108). ■

H.4 Proof for General k : attaining the lower bound

Given Lemma 23, we will be able to deduce the desired strict inclusion $\mathcal{C}_1 \supsetneq \mathcal{C}_2 \supsetneq \dots \supsetneq \mathcal{C}_k$ via the following result, which shows that the lower bound (107) or (108) is attainable and increasingly more restrictive:

Lemma 24 Fix $k \geq 2$ and suppose p_{ii} is sufficiently large, for $1 \leq i \leq 3$. Then $\lambda_j > (1 - p_{22} + p_{32})\eta_j > 0, \forall 1 \leq j \leq k$. Moreover there exists $\mu \in \mathcal{C}_k$ such that $\mu(1, 2) = \mu(3, 1) = 0$, and (107) or (108) holds with equality depending on whether p_{13} or p_{23} is larger.

Proof. As $p_{ii} \rightarrow 1$, it is clear from (103) that $\lambda_j \rightarrow 1$ and $\eta_j \rightarrow j - 1$ for $1 \leq j \leq k$. Thus we do have $\lambda_j > (1 - p_{22} + p_{32})\eta_j > 0$.

We will prove in detail that when $p_{13} \geq p_{23}$, the necessary condition (107) can be satisfied with equality. The opposite case with $p_{23} > p_{13}$ can be handled by almost the same arguments, which we omit.

Let us begin by investigating the conditions on $\nu(a^{-k}, \dots, a^{-1}, s, a)$ in order for (107) to hold. From the proof of Lemma 23, we see that the following ensures the equivalence between (106) and (107):

$$\begin{aligned} \nu(1, 3, 2) &= \nu(3, 3, 2) = 0. \\ \nu(a^{-j} = 2, 3, \dots, 3, s = 2, a = 1) &= 0, \quad \forall 1 \leq j \leq k. \end{aligned} \tag{109}$$

Since (106) comes from (100), we also impose the following to satisfy (100) with equality:

$$\nu(a^{-j} = 1, 3, \dots, 3, s = 1, a = 3) = 0, \quad \forall 1 \leq j \leq k \tag{110}$$

$$\nu(a^{-k} = \dots = a^{-1} = 3, s = 1, a = 3) = 0. \tag{111}$$

We will use these conditions to iteratively solve for $\nu(a^{-m}, \dots, a^{-1}, s, a)$, for $1 \leq m \leq k$. The starting point $\nu(a^{-1}, s, a)$ is straightforward:

Claim 22 Suppose p_{ii} is sufficiently large, then there exists $\nu(a^{-1}, s, a) \geq 0$ with $\nu(1, 1, 3) = \nu(1, 3, 2) = \nu(2, 2, 1) = \nu(3, 3, 2) = 0$ that solves the linear system for $\mu \in \mathcal{C}_1$.

Proof. Such a ν can be uniquely determined, and the results are mostly identical to those given in the proof of Claim 20. The difference here results from $\mu(1, 3) > \mu(3, 2)(p_{31} - p_{21})$, and so we only have a few entries changed:

$$\nu(2, 1, 3) = \mu(3, 2)(p_{31} - p_{21}).$$

$$\nu(3, 1, 3) = \mu(1, 3) - \nu(2, 1, 3) = \mu(1, 3) - \mu(3, 2)(p_{31} - p_{21}).$$

$$\begin{aligned}\nu(3, 2, 3) &= \mu(1, 3)(p_{33} - p_{13}) + \mu(2, 3)(p_{33} - p_{23}) - \mu(1, 3) + \mu(3, 2)(p_{31} - p_{21}) \\ &= \mu(3, 2)(p_{22} - p_{32}) - \mu(1, 3)(1 + p_{13} - p_{23}).\end{aligned}$$

Since μ satisfies (107) with equality, we have $\frac{\mu(1,3)}{\mu(3,2)} \rightarrow 0$ as $p_{ii} \rightarrow 1$. Thus $\nu(3, 2, 3) > 0$, completing the proof. ■

The induction step is summarized by the following claim:

Claim 23 *Suppose that, for some $2 \leq m \leq k$, we have found $\nu(a^{-(m-1)}, \dots, a^{-1}, s, a) \geq 0$ that solves the linear system for \mathcal{M}_{m-1} , obeys (109) to (111) and satisfies the following proportionality condition:*

$$\frac{\nu(a^{-(m-1)}, \dots, a^{-2}, a^{-1} = 2, s = 1, a = 3)}{\nu(a^{-(m-1)}, \dots, a^{-2}, a^{-1} = 2, s = 2, a = 3)} = \frac{\nu(2, 1, 3)}{\nu(2, 2, 3)}, \forall a^{-(m-1)}, \dots, a^{-2}. \quad (112)$$

Then this ν can be extended to $\nu(a^{-m}, \dots, a^{-1}, s, a) \geq 0$ that solves the linear system for \mathcal{M}_m , obeys (109) to (111) and satisfies the corresponding proportionality condition for all a^{-m}, \dots, a^{-2} .

Proof. By our previous discussion of the iterated linear system, we need to solve for $\nu(a^{-m}, \dots, a^{-1}, s, a)$ from the following system of equations:

$$\begin{aligned}\sum_{a^{-m}} \nu(a^{-m}, \dots, a^{-1}, s, a) &= \nu(a^{-(m-1)}, \dots, a^{-1}, s, a). \\ \sum_s \nu(a^{-m}, \dots, a^{-1}, s, a) &= \lambda(a^{-m})p(a^{-(m-1)}|a^{-m}) \cdots p(a|a^{-1}). \\ \sum_a \nu(a^{-m}, \dots, a^{-1}, s, a) &= \sum_{s^{-1}} \nu(a^{-m}, \dots, a^{-2}, s^{-1}, a^{-1})p(s|s^{-1}).\end{aligned}$$

We take $a^{-(m-1)}, \dots, a^{-1}$ as fixed and consider the above as a $3 \times 3 \times 3$ transportation problem. First consider $a^{-1} = 1$. Because $\nu(1, 1, 2) = \nu(1, 1, 3) = \nu(1, 3, 1) = \nu(1, 3, 2) = 0$, we must have $\nu(a^{-m}, \dots, 1, 1, 2) = \cdots = 0$. Then we can uniquely solve for $\nu(a^{-m}, \dots, a^{-1} = 1, s, a)$ from the required row and column sums. The only off-diagonal entry remaining is:

$$\nu(a^{-m}, \dots, a^{-1} = 1, s = 2, a = 1) = \nu(a^{-m}, \dots, a^{-1} = 2, s^{-1} = 2, a^{-1} = 1)(p_{11} - p_{21}).$$

This is positive, so we are done for the case $a^{-1} = 1$ (as before, diagonal entries being negative is not an issue because we are concerned with the cone \mathcal{C}_1).

Next consider $a^{-1} = 2$. Here we have $\nu(2, 1, 2) = \nu(2, 2, 1) = \nu(2, 3, 1) = 0$, from which we obtain:

$$\nu(a^{-m}, \dots, a^{-1} = 2, s = 1, a = 3) = \nu(a^{-m}, \dots, a^{-2}, s^{-1} = 3, a^{-1} = 2)(p_{31} - p_{21}).$$

From the proportionality condition we deduce:

$$\nu(a^{-m}, \dots, a^{-1} = 2, s = 2, a = 3) = \nu(a^{-m}, \dots, a^{-2}, s^{-1} = 3, a^{-1} = 2)(p_{31} - p_{21}) \frac{\nu(2, 2, 3)}{\nu(2, 1, 3)}.$$

The remaining off-diagonal entry is thus given by:

$$\begin{aligned} & \nu(a^{-m}, \dots, a^{-1} = 2, s = 3, a = 2) - \nu(a^{-m}, \dots, a^{-1} = 2, s = 2, a = 3) \\ &= \sum_s \nu(a^{-m}, \dots, a^{-1} = 2, s, a = 2) - \sum_a \nu(a^{-m}, \dots, a^{-1} = 2, s = 2, a) \\ &= \nu(a^{-m}, \dots, a^{-2}, s^{-1} = 3, a^{-1} = 2)(p_{22} - p_{32}). \end{aligned}$$

Therefore $\nu(a^{-m}, \dots, a^{-1} = 2, s = 2, a = 3)$ and $\nu(a^{-m}, \dots, a^{-1} = 2, s = 3, a = 2)$ are both positive as desired.

Lastly, consider $a^{-1} = 3$, where we know $\nu(3, 1, 2) = \nu(3, 3, 1) = \nu(3, 3, 2) = 0$. If a^{-m}, \dots, a^{-2} are not all equal to state ‘‘3’’, let $2 \leq t \leq m$ be the smallest index such that $a^{-t} \neq 3$. We distinguish three possibilities:

1. $a^{-t} = 1$. From (110) we have $\nu(a^{-m}, \dots, a^{-1} = 3, s = 1, a = 3) = 0$. We can then solve the other entries of ν :

$$\begin{aligned} & \nu(a^{-m}, \dots, a^{-1} = 3, s = 2, a = 1) \\ &= \sum_s \nu(a^{-m}, \dots, a^{-1} = 3, s, a = 1) - \sum_a \nu(a^{-m}, \dots, a^{-1} = 3, s = 1, a) \quad (113) \\ &= \nu(a^{-m}, \dots, a^{-2}, s^{-1} = 2, a^{-1} = 3)(p_{31} - p_{21}). \end{aligned}$$

In the last step, we used $\nu(a^{-m}, \dots, a^{-2}, s^{-1} = 1, a^{-1} = 3) = 0$, which is due to $a^{-t} = 1, a^{-(t-1)} = \dots = a^{-2} = 3$ and the induction hypothesis (110). Similarly:

$$\nu(a^{-m}, \dots, a^{-1} = 3, s = 2, a = 3) = \nu(a^{-m}, \dots, a^{-2}, s^{-1} = 2, a^{-1} = 3)(p_{33} - p_{23}). \quad (114)$$

These are both positive, as desired.

2. $a^{-t} = 2$. From (109) we know that $\nu(a^{-m}, \dots, a^{-1} = 3, s = 2, a = 1) = 0$. The remaining off-diagonal entries $\nu(\dots, 3, 1, 3)$ and $\nu(\dots, 3, 2, 3)$ can then be solved:

$$\begin{aligned}
& \nu(a^{-m}, \dots, a^{-1} = 3, s = 1, a = 3) \\
&= \sum_a \nu(a^{-m}, \dots, a^{-1} = 3, s = 1, a) - \sum_s \nu(a^{-m}, \dots, a^{-1} = 3, s, a = 1) \\
&= \nu(a^{-m}, \dots, a^{-2}, s^{-1} = 1, a^{-1} = 3)(p_{11} - p_{31}) - \nu(a^{-m}, \dots, a^{-2}, s^{-1} = 2, a^{-1} = 3)(p_{31} - p_{21}).
\end{aligned} \tag{115}$$

To show this is positive, let us note that the proportionality condition (112) generalizes to the following:

$$\frac{\nu(a^{-m}, \dots, a^{-t} = 2, 3, \dots, 3, s^{-1} = 1, a^{-1} = 3)}{\nu(a^{-m}, \dots, a^{-t} = 2, 3, \dots, 3, s^{-1} = 2, a^{-1} = 3)} = \frac{\nu(a^{-t} = 2, 3, \dots, 3, s^{-1} = 1, a^{-1} = 3)}{\nu(a^{-t} = 2, 3, \dots, 3, s^{-1} = 2, a^{-1} = 3)}. \tag{116}$$

The case $t = 2$ is exactly the proportionality condition assumed above, and the general result follows by exploiting recursive formulae of the form (115).

Given this generalized proportionality condition, the R.H.S. of (115) is positive if $\nu(a^{-t} = 2, 3, \dots, 3, s^{-1} = 1, a^{-1} = 3)(p_{11} - p_{31}) \geq \nu(a^{-t} = 2, 3, \dots, 3, s^{-1} = 2, a^{-1} = 3)(p_{31} - p_{21})$. Plugging in the formulae (104) and (105) as well as conditions (109) and (110),⁴⁹ we simply need to show $\nu(2, 1, 3)\lambda_t \geq \nu(2, 2, 3)(p_{31} - p_{21})\eta_t$. We have computed that $\nu(2, 1, 3) = \mu(3, 2)(p_{31} - p_{21})$ and $\nu(2, 2, 3) = \mu(1 - p_{22} + p_{32})$. Thus, the desired inequality reduces to $\lambda_t \geq (1 - p_{22} + p_{32})\eta_t$, which we have shown earlier.

In a similar way we have:

$$\begin{aligned}
& \nu(a^{-m}, \dots, a^{-1} = 3, s = 2, a = 3) \\
&= \sum_a \nu(a^{-m}, \dots, a^{-1} = 3, s = 2, a) - \sum_s \nu(a^{-m}, \dots, a^{-1} = 3, s, a = 2) \\
&= \nu(a^{-m}, \dots, a^{-2}, s^{-1} = 1, a^{-1} = 3)(p_{12} - p_{32}) + \nu(a^{-m}, \dots, a^{-2}, s^{-1} = 2, a^{-1} = 3)(p_{22} - p_{32}).
\end{aligned} \tag{117}$$

⁴⁹These are all valid since $\nu(a^{-t}, \dots, a^{-2}, s^{-1}, a^{-1})$ solves the linear system for \mathcal{M}_{t-1} by the induction hypothesis.

Using (116) and plugging in the formulae (104) and (105) again, we conclude that the R.H.S. above is positive if $\nu(2, 1, 3)(p_{32} - p_{12})\eta_t \leq \nu(2, 2, 3)(\lambda_t + (p_{22} - p_{32} - p_{11} + p_{31})\eta_t)$. This further simplifies to $(p_{31} - p_{21})(p_{32} - p_{12})\eta_t \leq (1 - p_{22} + p_{32})(\lambda_t + (p_{22} - p_{32} - p_{11} + p_{31})\eta_t)$, which holds because $0 < p_{31} - p_{21} \leq 1 - p_{11} + p_{31}$, $p_{32} - p_{12} < 1 - p_{22} + p_{32}$ and $\lambda_t \geq (1 - p_{22} + p_{32})\eta_t$. This completes the proof that ν is positive off the diagonal.

3. $a^{-m} = \dots = a^{-1} = 3$. Here the off-diagonal entries are $\nu(\dots, 3, 1, 2) = \nu(\dots, 3, 3, 1) = \nu(\dots, 3, 3, 2) = 0$ and $\nu(\dots, 3, 1, 3), \nu(\dots, 3, 2, 1), \nu(\dots, 3, 2, 3)$. We compute using (110) and (104) that:

$$\begin{aligned}
& \nu(a^{-m} = \dots = a^{-1} = 3, s = 1, a = 3) \\
&= \mu(1, 3) - \sum_{j=1}^m \nu(a^{-j} = 2, a^{-(j-1)} = \dots = a^{-1} = 3, s = 1, a = 3) \\
&= \mu(1, 3) - \mu(3, 2)(p_{31} - p_{21}) \sum_{j=1}^m (\lambda_j - (1 - p_{22} + p_{32})\eta_j).
\end{aligned} \tag{118}$$

This is positive by the assumption that (107) holds equal and $\lambda_j > (1 - p_{22} + p_{32})\eta_j$. In particular, at $m = k$ we have $\nu(a^{-k} = \dots = a^{-1} = 3, s = 1, a = 3)$, as required by (111).

From (109) we similarly obtain:

$$\begin{aligned}
& \nu(a^{-m} = \dots = a^{-1} = 3, s = 2, a = 1) \\
&= \nu(3, 2, 1) - \sum_{j=2}^m \nu(a^{-j} = 1, a^{-(j-1)} = \dots = a^{-1} = 3, s = 2, a = 1) \\
&= \mu(1, 3)(1 - p_{11} + p_{21}) - \sum_{j=1}^{m-1} \nu(a^{-j} = 1, a^{-(j-1)} = \dots = a^{-1} = 3, s = 2, a = 3)(p_{31} - p_{21}) \\
&= \mu(1, 3)(1 - p_{11} + p_{21}) - \mu(1, 3) \cdot (p_{13} - p_{23})(p_{31} - p_{21}) \sum_{j=1}^{m-1} (p_{33} - p_{23})^{j-1}.
\end{aligned} \tag{119}$$

In the last two steps above we used recursive formulae of the form (113) and (114) and $\nu(1, 2, 3) = \mu(1, 3)(p_{13} - p_{23})$. Since $1 - p_{11} + p_{21} > p_{13} - p_{23}$ and $(p_{31} - p_{21}) \sum_{j=1}^{m-1} (p_{33} - p_{23})^{j-1} < \frac{p_{31} - p_{21}}{1 - p_{33} + p_{23}} < 1$, this ν is always positive.

Finally we compute that:

$$\begin{aligned}
& \nu(a^{-m} = \dots = a^{-1} = 3, s = 2, a = 3) \\
&= \mu(2, 3) - \sum_{j=1}^m \nu(a^{-j} = 1, 3, \dots, 3, s = 2, a = 3) - \sum_{j=1}^m \nu(a^{-j} = 2, 3, \dots, 3, s = 2, a = 3) \\
&= (\mu(3, 2) - \mu(1, 3)) - \mu(1, 3)(p_{13} - p_{23}) \frac{1 - (p_{33} - p_{23})^m}{1 - p_{33} + p_{23}} \\
&\quad - \mu(3, 2)(1 - p_{22} + p_{32}) \sum_{j=1}^m \lambda_j \\
&\quad - \mu(3, 2) ((p_{31} - p_{21})(p_{12} - p_{32}) + (1 - p_{22} + p_{32})(p_{22} - p_{32} - p_{11} + p_{31})) \sum_{j=1}^m \eta_j.
\end{aligned} \tag{120}$$

In the last step we used the calculation just now for $\sum_{j=1}^m \nu(a^{-j} = 1, 3, \dots, 3, s = 2, a = 3)$ as well as formula (105). Now recall that as $p_{ii} \rightarrow 1$, $\lambda_j \rightarrow 1$ while $\eta_j \rightarrow j-1$. Thus $(1 - p_{22} + p_{32}) \sum_{j=1}^m \lambda_j$ and $((p_{31} - p_{21})(p_{12} - p_{32}) + (1 - p_{22} + p_{32})(p_{22} - p_{32} - p_{11} + p_{31})) \sum_{j=1}^m \eta_j$ both vanish. Furthermore, since we assume (107) holds equal,

$$\frac{\mu(1, 3)}{\mu(3, 2)(1 - p_{33} + p_{23})} < \frac{\mu(1, 3)}{\mu(3, 2)(p_{31} - p_{21})} < \sum_{j=1}^k \lambda_j$$

remains bounded. These observations together imply that the R.H.S. of (120) is positive as $p_{ii} \rightarrow 1$.

With that we have verified $\nu(a^{-m}, \dots, a^{-1}, s, a)$ is indeed positive. This completes the proof of the claim. ■

Lemma 24 follows, and so does Proposition 7'. ■

H.5 Zero Transitions

In this subsection, we provide an example to show that the full-support assumption stated in Proposition ?? is necessary for the result there. Consider the following cyclic transition matrix:

$$P = \begin{pmatrix} 0.4 & 0.6 & 0 \\ 0 & 0.4 & 0.6 \\ 0.6 & 0 & 0.4 \end{pmatrix}$$

In \mathcal{M}_1 , a vertex adjacent to μ^{tt} is

$$\mu(s, a) = \begin{pmatrix} \frac{5}{24} & 0 & \frac{1}{8} \\ \frac{1}{8} & \frac{1}{8} & \frac{1}{12} \\ 0 & \frac{5}{24} & \frac{1}{8} \end{pmatrix}$$

The corresponding $\nu(a^{-1}, s, a)$ has the following configuration:

$$\nu = \begin{pmatrix} + & 0 & 0 \\ + & + & 0 \\ 0 & + & 0 \end{pmatrix} \begin{pmatrix} 0 & 0 & + \\ 0 & + & + \\ 0 & + & + \end{pmatrix} \begin{pmatrix} + & 0 & 0 \\ + & 0 & + \\ 0 & 0 & + \end{pmatrix}$$

Note that this configuration nests under sub-case (B7) in our proof of Proposition ?? before.

One routinely verifies that for any $\varepsilon \in (0, 1)$, $\nu_\varepsilon = \varepsilon\nu + (1 - \varepsilon)\nu^{tt}$ is the unique invariant distribution associated with $\mu_\varepsilon = \varepsilon\mu + (1 - \varepsilon)\mu^{tt}$. For μ_ε to be in \mathcal{M}_2 , ν_ε must be extended to some $\hat{\nu}(a^{-2}, a^{-1}, s, a)$ that satisfies the iterative linear system. Since $\hat{\nu}(3, 3, 3, 1) = \nu_\varepsilon(3, 3, 1) = 0$ and $\hat{\nu}(3, 3, 3, 2) = \nu_\varepsilon(3, 3, 2) = 0$, we have

$$\begin{aligned} \hat{\nu}(3, 3, 3, 3) &= \sum_a \hat{\nu}(3, 3, 3, a) \\ &= \nu_\varepsilon(3, 1, 3)p_{13} + \nu_\varepsilon(3, 2, 3)p_{23} + \nu_\varepsilon(3, 3, 3)p_{33} \\ &= \nu_\varepsilon(3, 2, 3)p_{23} + \nu_\varepsilon(3, 3, 3)p_{33}. \end{aligned}$$

Thus from $\hat{\nu}(3, 3, 1, 3) = \nu_\varepsilon(3, 1, 3) = 0$ we deduce

$$\begin{aligned}\hat{\nu}(3, 3, 2, 3) &= \sum_s \hat{\nu}(3, 3, s, 3) - \hat{\nu}(3, 3, 3, 3) \\ &= (\nu_\varepsilon(3, 1, 3) + \nu_\varepsilon(3, 2, 3) + \nu_\varepsilon(3, 3, 3))p_{33} - \hat{\nu}(3, 3, 3, 3) \\ &= \nu_\varepsilon(3, 2, 3)(p_{33} - p_{23}) < 0.\end{aligned}$$

This contradiction shows that $\mu \notin \mathcal{C}_2$, so that $\mathcal{C}_1 \neq \mathcal{C}_2$ for this process despite all transition probabilities being below $\frac{\sqrt{5}-1}{2}$.

What goes wrong if we follow the proof of sub-case (B7) and set $\tau(s, a|s^{-1}, a^{-1})$ according to Eq. (68)? The issue is that for such a τ , $\nu(a^{-2} = 3, a^{-1} = 3, s = 2, a = 2)$ will be negative no matter how small we choose ε . Because $\nu^{tt}(a^{-2} = 3, a^{-1} = 3, s = 2, a = 2) = 0$ due to $p_{32} = 0$, Lemma 20 fails.

I Recursive representations

With the notations of Section 4.3, we here prove that $\sup_{L,M} S^{k,L,M} < 1$, for each k .

The intuition for the proof is straightforward, as discussed in the body of the paper. The only challenge is the following. To get a payoff arbitrarily close to 1 in the finitely repeated game, the equilibrium strategy of player 1 must be close to truth-telling (in a sense made precise below), but it need not be exactly truth-telling. The strategy that implements μ^{tt} (the candidate for a profitable deviation by player 1) is indistinguishable from exact truth-telling, but it need not be indistinguishable from a strategy close to it. Hence, one must show that for every strategy close to truth-telling, there is a strategy close to the one implementing μ^{tt} that is indistinguishable from it. That this is possible is a consequence of Lemma 12 below, which builds itself on Lemma 11.

Lemma 11 *Let T and B be finite sets, and $\mu \in \Delta(T \times B)$. For each $\tilde{\mu}|_B \in \Delta(B)$, there exists a distribution $\hat{\mu} \in \Delta(T \times B)$ such that*

$$\hat{\mu}|_T = \mu|_T, \hat{\mu}|_B = \tilde{\mu}|_B \text{ and } \|\mu - \hat{\mu}\|_1 = \|\mu|_B - \tilde{\mu}|_B\|_1.$$

In this statement, $\mu|_T$ stands for the marginal of μ on T . Lemma 11 is a finite-dimensional version of Lemma 8.2 in Laraki (2004). For completeness, we provide a proof.

Proof. The distribution $\hat{\mu}$ will be obtained as the solution to a linear program. Define $B_+ = \{b, \tilde{\mu}|_B(b) \geq \mu|_B(b)\}$ and $B_- = \{b, \tilde{\mu}|_B(b) \leq \mu|_B(b)\}$. Elements of B_+ (resp., B_-) should see their weight increase (resp., decrease). Consider the linear program

$$\sup \left(\sum_{t \in T, b \in B_+} \delta_{tb} - \sum_{t \in T, b \in B_-} \delta_{tb} \right),$$

subject to the constraints:

C1 $\delta_{tb} \geq 0$ for all $b \in B_+$ and $\delta_{tb} \leq 0$ for all $b \in B_-$, all $t \in T$.

C2 $0 \leq \mu(t, b) + \delta_{tb} \leq 1$, all $t \in T$, $b \in B$.

C3 $\sum_{b \in B} \delta_{tb} = 0$, all $t \in T$.

C4 $\mu|_B(b) + \sum_{t \in T} \delta_{tb} \leq \tilde{\mu}|_B(b)$, for all $b \in B_+$, and $\mu|_B(b) + \sum_{t \in T} \delta_{tb} \geq \tilde{\mu}|_B(b)$ for all $b \in B_-$.

This problem is feasible ($\delta = 0$) and bounded (by **C2**), hence has an optimal solution $\hat{\delta}$. We will show that $\hat{\mu}(t, b) := \mu(t, b) + \hat{\delta}_{tb}$ is the desired distribution.

Conditions **C2** and **C3** ensure that $\hat{\mu} \in \Delta(T \times B)$, with $\hat{\mu}|_T = \mu|_T$. Condition **C1** implies that $\|\hat{\mu} - \mu\|_1 = \sum_{t, b \in B_+} \hat{\delta}_{tb} - \sum_{t, b \in B_-} \hat{\delta}_{tb}$. Condition **C4** implies that $\hat{\mu}|_B(b) \leq \tilde{\mu}|_B(b)$ if and only if $b \in B_+$. Hence, the value of the program cannot exceed $\|\mu|_B - \tilde{\mu}|_B\|_1$ and equality holds iff $\hat{\mu}|_B = \tilde{\mu}|_B$. In that case, it follows that $\|\mu - \hat{\mu}\|_1 = \|\mu|_B - \tilde{\mu}|_B\|_1$.

Thus, it suffices to check that $\hat{\mu}|_B = \tilde{\mu}|_B$. Assume to the contrary that $\hat{\mu}|_B(b) < \tilde{\mu}|_B(b)$ and $\hat{\mu}|_B(b') > \tilde{\mu}|_B(b')$ for some $b \in B_+$ and $b' \in B_-$. Choose t such that $\hat{\mu}(t, b') > 0$ –so that $\hat{\mu}(t, b) < 1$. Increasing slightly $\hat{\mu}(t, b)$ at the expense of $\hat{\mu}(t, b')$ improves upon $\hat{\mu}$ –a contradiction.

■

Lemma 12 *Let σ_1 be a strategy of player 1 in $\Gamma^k(t)$, and let $\tilde{\lambda}_k \in \Delta(A^k)$ be arbitrary. Then there exists a strategy $\hat{\sigma}_1$ such that*

$$\mathbf{P}_{\lambda, \hat{\sigma}_1}|_{A^k} = \tilde{\lambda}_k \text{ and } \|\mathbf{P}_{\lambda, \hat{\sigma}_1} - \mathbf{P}_{\lambda, \sigma_1}\|_1 \leq C_k \|\mathbf{P}_{\lambda, \sigma_1} - \tilde{\lambda}_k\|_1,$$

for some constant C_k .

In this statement, $\mathbf{P}_{\lambda, \sigma_1}$ is the distribution of the sequence of player 1's actions induced by σ_1 . For notational simplicity, we also abstract from the actions of player 2 and 3, whose strategies are fixed.

Proof. We use induction over k . For $k = 1$, the claim coincides with the conclusion of Lemma 11. Let now a strategy σ_1 in $\Gamma^{k+1}(t)$ and $\tilde{\lambda}_{k+1} \in \Delta(A^{k+1})$ be given, and denote by $\sigma_{1,k}$ and $\tilde{\lambda}_k$ their restrictions to the first k rounds. Using the induction hypothesis, let $\hat{\sigma}_{1,k}$ be a strategy over the first k rounds such that $\mathbf{P}_{\lambda, \sigma_{1,k}}|_{A^k} = \tilde{\lambda}_k$ and $\|\mathbf{P}_{\lambda, \sigma_{1,k}} - \mathbf{P}_{\lambda, \hat{\sigma}_{1,k}}\|_1 \leq C_k \|\mathbf{P}_{\lambda, \sigma_{1,k}}|_{A^k} - \tilde{\lambda}_k\|_1$.

Let $\bar{\sigma}_1$ be the strategy in $\Gamma^{k+1}(t)$ which coincides with $\hat{\sigma}_{1,k}$ in the first k rounds, and with σ_1 in the last one.⁵⁰ Fix $\vec{a} \in A^k$. Applying Lemma 11 with $T = S^{k+1}$, $B = A$, $\mu := \mathbf{P}_{\lambda, \bar{\sigma}_1}(\cdot | \vec{a})$ and $\tilde{\mu}|_A = \tilde{\lambda}_{k+1}(\cdot | \vec{a}) \in \Delta(A)$, we get $\hat{\mu}_{\vec{a}} \in \Delta(T \times B)$ with the distributional properties stated there.

We then define $\hat{\sigma}_1$ in round $k + 1$ by setting $\hat{\sigma}_1(\cdot | \vec{a}, \vec{s}) := \mu_{\vec{a}}(\cdot | \vec{s}) \in \Delta(A)$ for each $\vec{s} \in S^{k+1} = T$. By construction and since $\mathbf{P}_{\lambda, \hat{\sigma}_1}$ coincides on A^k with $\tilde{\lambda}_k$, the marginal of $\mathbf{P}_{\lambda, \hat{\sigma}_1}$ on A^{k+1} is equal to $\tilde{\lambda}_{k+1}$.

In addition, because σ_1 and $\bar{\sigma}_1$ coincide in round $k + 1$ and by the induction hypothesis, one has

$$\|\mathbf{P}_{\lambda, \sigma_1} - \mathbf{P}_{\lambda, \bar{\sigma}_1}\|_1 = \|\mathbf{P}_{\lambda, \sigma_1}|_{(S \times A)^k} - \mathbf{P}_{\lambda, \bar{\sigma}_1}|_{(S \times A)^k}\|_1 \leq C_k \|\mathbf{P}_{\lambda, \sigma_1}|_{A^k} - \tilde{\lambda}_k\|_1. \quad (121)$$

On the other hand,

$$\begin{aligned} \|\mathbf{P}_{\lambda, \hat{\sigma}_1} - \mathbf{P}_{\lambda, \bar{\sigma}_1}\|_1 &= \sum_{\vec{a}} \mathbf{P}_{\lambda, \bar{\sigma}_1}(\vec{a}) \times \sum_{a \in A, \vec{s} \in S^{k+1}} |\mathbf{P}_{\lambda, \hat{\sigma}_1}(a, \vec{s} | \vec{a}) - \mathbf{P}_{\lambda, \bar{\sigma}_1}(a, \vec{s} | \vec{a})| \\ &= \sum_{\vec{a}} \mathbf{P}_{\lambda, \bar{\sigma}_1}(\vec{a}) \times \|\hat{\mu}_{\vec{a}} - \mu\|_1 \\ &= \sum_{\vec{a}} \mathbf{P}_{\lambda, \bar{\sigma}_1}(\vec{a}) \times \|\mu|_A - \tilde{\lambda}_{k+1}(\cdot | \vec{a})\|_1 \\ &= \|\mathbf{P}_{\lambda, \bar{\sigma}_1}|_{A^{k+1}} - \tilde{\lambda}_{k+1}\|_1 \\ &\leq \|\mathbf{P}_{\lambda, \sigma_1}|_{A^{k+1}} - \mathbf{P}_{\lambda, \bar{\sigma}_1}|_{A^{k+1}}\|_1 + \|\mathbf{P}_{\lambda, \sigma_1}|_{A^{k+1}} - \tilde{\lambda}_{k+1}\|_1 \\ &= \|\mathbf{P}_{\lambda, \sigma_1}|_{A^k} - \mathbf{P}_{\lambda, \bar{\sigma}_1}|_{A^k}\|_1 + \|\mathbf{P}_{\lambda, \sigma_1}|_{A^{k+1}} - \tilde{\lambda}_{k+1}\|_1 \\ &\leq (C_k + 1) \|\mathbf{P}_{\lambda, \sigma_1}|_{A^{k+1}} - \tilde{\lambda}_{k+1}\|_1, \end{aligned}$$

where the first equality is an identity, the second one follows from the definition of $\hat{\sigma}_1$ in round $k + 1$, the third one follows since $\hat{\mu}_{\vec{a}}$ satisfies the conclusion of Lemma 11, the first inequality follows from the triangle inequality, and the rest follows from the same lines as (121).

⁵⁰In particular, the distribution of the first k actions of player 1 is $\tilde{\lambda}_k$.

Adding the last inequality and (121) gives the result, with $C_{k+1} = 2C_k + 1$. ■

We proceed to the proof of $\sup_{L \geq 1} S^{k,L,M} < 1$. Assume to the contrary that for each $\varepsilon > 0$, there is $L_\varepsilon \geq 1$, t_ε , and an equilibrium $\sigma_\varepsilon \in E^{k,L_\varepsilon,M}(t_\varepsilon)$ such that $\sum_i v^i(\sigma_\varepsilon) \geq 1 - \varepsilon$. Since total payoffs equal to 1 can only be obtained when player 1 reports truthfully, this implies that $\mathbf{P}_{\lambda,\sigma_\varepsilon}$ converges as $\varepsilon \rightarrow 0$, to the distribution on $(S \times A)^k$ induced by truth-telling. This implies in turn that the equilibrium payoff vector $(v^i(\sigma_\varepsilon))_i$ —net of transfers—converges to $\mu^{tt} \cdot r$.

Now, fix a stationary reporting policy σ_* and an invariant distribution ν_* implementing μ_{x_k} . The reporting policy σ_* can be viewed as a strategy σ_1 of player in $\Gamma^k(t_\varepsilon)$, when letting player 1 draw in round 1 a fictitious past according to $\nu_*(\cdot | s_1)$. In particular, the expected payoff vector, net of transfers, induced by $(\sigma_1, \sigma_{-1,\varepsilon})$ is equal to $\mu_{x_k} \cdot r$.

Apply next Lemma 12 with σ_1 and $\tilde{\lambda} := \mathbf{P}_{\lambda,\sigma_\varepsilon}|_{A^k}$ to get a new strategy $\hat{\sigma}_{1,\varepsilon}$. Since $\mathbf{P}_{\lambda,\hat{\sigma}_{1,\varepsilon}}|_{A^k} = \mathbf{P}_{\lambda,\sigma_{1,\varepsilon}}|_{A^k}$, the expected transfers to player 1 do not change when deviating from $\sigma_{1,\varepsilon}$ to $\hat{\sigma}_{1,\varepsilon}$. Since $\|\mathbf{P}_{\lambda,\sigma_1} - \mathbf{P}_{\lambda,\hat{\sigma}_{1,\varepsilon}}\|_1 \leq C_k \|\mathbf{P}_{\lambda,\sigma_{1,\varepsilon}}|_{A^k} - \lambda\|_1$ converges to zero as $\varepsilon \rightarrow 0$, expected payoffs net of transfers under $(\hat{\sigma}_{1,\varepsilon}, \sigma_{-1,\varepsilon})$ converge to $\mu_{x_k} \cdot r$.

Since $\mu_{x_k} \cdot r^1 > \mu^{tt} \cdot r^1$, it follows that $\hat{\sigma}_{1,\varepsilon}$ is a profitable deviation upon $\sigma_{1,\varepsilon}$ for ε small enough.

References

- [1] Laraki, R. (2004). “On the Regularity of the Convexification Operator on a Compact Set,” *Journal of Convex Analysis*, **11**, 209–234.
- [2] Smith, G. and J. Dawson (1979). “The determination of necessary and sufficient conditions for the existence of a solution to the $3 \times 3 \times 3$ multi-index problem,” *Aplikace Matematiky*, **24**, 201–208. 1/CRC.