Non-Recursive Dynamic Incentives: A Rate of Convergence Approach*

Takuo Sugaya — Stanford GSB
Alexander Wolitzky — MIT

July 18, 2024

Abstract

In repeated principal-agent problems and games, more outcomes are implementable when performance signals are privately observed by a principal or mediator with commitment power than when the same signals are publicly observed and form the basis of a recursive equilibrium. We investigate the gains from non-recursive equilibria (e.g., "review strategies") based on privately observed signals. Under a pairwise identifiability condition, we find that the gains from non-recursive equilibria are "small": their inefficiency is of the same $1-\delta$ power order as that of recursive equilibria. Thus, while private strategies or monitoring can outperform public ones for a fixed discount factor, they cannot accelerate the power rate of convergence to the efficient payoff frontier when the folk theorem holds. An implication is that the gains from withholding performance feedback from agents are small when the parties are patient.

Keywords: repeated games, repeated agency, imperfect monitoring, performance feedback, review strategies, rate of convergence, folk theorem, martingales

JEL codes: C72, C73

*Previous titles: "Performance Feedback in Long-Run Relationships: A Rate of Convergence Approach," "Rate of Convergence in Repeated Games: A Universal Speed Limit." For helpful comments, we thank seminar participants at Kyoto, Stanford, and Warwick, as well as Drew Fudenberg, Johannes Hörner, Stephen Morris, Yuliy Sannikov, Andrzej Skrzypacz, and Satoru Takahashi.

1 Introduction

Most analysis of repeated moral hazard problems and games focuses on contracts and equilibria that are recursive in the players' continuation values. This approach is without loss in single-agent problems with public performance signals (Spear and Srivastava, 1987). It is also without loss in repeated games with imperfect public monitoring, if attention is restricted to equilibria in pure strategies or in strategies that depend only on the public signals (Abreu, Pearce, and Stacchetti, 1990; Fudenberg, Levine, and Maskin, 1994). In contrast, in single-agent problems where the principal privately observes performance, or in repeated games where signals are privately observed by a mediator, more payoffs are implementable as compared to the case where the same signals are publicly observed, as concealing signals reduces the players' available deviations. Similarly, focusing on pure or public equilibria in repeated games with public monitoring is not without loss (Kandori and Obara, 2006). Yet, characterizing the equilibrium payoff set with private signals or strategies is intractable, precisely because this set lacks a tractable recursive structure (Kandori, 2002). The extent of the possible gains from non-recursive equilibria based on private signals or strategies over recursive equilibria based on public signals is thus an open question, which forms the subject of the current paper.

Specifically, we consider discounted repeated games where in each period players take actions a and a signal y is drawn from a distribution $p(y|a)$ with non-moving support. We compare the equilibrium payoff sets in a version of the game with public monitoring, where the signal y is publicly observed and attention is restricted to equilibria in public strategies, and a version with private monitoring, where the signal y is observed only by a principal or mediator with commitment power, who privately recommends actions to the players. We call these two versions of the game the *public game* and the *blind game*. By the revelation principle, for any discount factor δ, the equilibrium payoff set is weakly larger in the blind game than the public game. Our question is, how much larger?

For any fixed discount factor $\delta < 1$, this question is difficult to answer in any generality, because characterizing equilibrium payoffs in the blind game is intractable. We instead adopt a *rate of convergence* approach: under standard identification conditions that ensure that efficiency is attainable in the $\delta \to 1$ limit, how quickly does inefficiency vanish as $\delta \to 1$ in the most efficient equilibrium in the public game as compared to the blind game?

Our main result is that inefficiency is of the same power order of $1 - \delta$ in both games. Thus, while private strategies or monitoring can outperform public ones for a fixed discount factor, they cannot accelerate the power rate of convergence to the efficient payoff frontier when the folk theorem holds. In this sense, the gains from non-recursive equilibria are small.

Our results have implications for the design of principal-agent relationships. An important design variable in such relationships is the amount of performance feedback provided to the agent. While providing feedback can have practical benefits that are not captured by our model, a benefit of *withholding* feedback is that this facilitates non-recursive contracting by making the game blind rather than public. However, our results show that this benefit of withholding feedback is small when the parties are patient.

The high-level intuition for our results is that, as compared to a recursive contract where the agents' continuation values are revealed in every period, pooling information across periods improves monitoring precision but also necessitates larger rewards and punishments, which reduces the scope for providing incentives by transferring surplus over time rather than destroying it. Our analysis shows that these two effects essentially cancel out, so that little is gained by pooling information.

A subtlety in our results is that, while inefficiency is always of the same power order in the public game and the blind game, this order depends on the curvature of the boundary of the feasible payoff set. If the boundary is smooth with positive curvature (as in Green and Porter, 1984, Spear and Srivastava, 1987, Sannikov, 2007, 2008, or Sadzik and Stacchetti, 2015), inefficiency is of order $1 - \delta$.[1] In this case, the first-order inefficiency associated with small continuation payoff movements along the payoff boundary is zero. We show that this implies that inefficiency in the public and blind games differs only by a constant factor: i.e., the rate of convergence is identical. Moreover, for a class of smooth principal-agent models (similar to Spear and Srivastava, 1987, or Sannikov, 2008), inefficiency in the public and blind games is identical up to a first-order approximation.

[1] For a class of continuous-time principal-agent problems with public monitoring, this was already observed by Sannikov (2008).

In contrast, if the boundary of the feasible payoff set is kinked (as in the case with discrete actions), inefficiency is of power order $(1-\delta)^{1/2}$.[2] In this case, the first-order inefficiency associated with small continuation payoff movements is positive. We show that this greater inefficiency leads to a greater value of withholding feedback: now, inefficiency in the public and blind games can differ by a log factor in $1-\delta$. Thus, while the value of withholding feedback is always "small" (no improvement in the power rate of convergence), it is somewhat less small in the kinked case (where there can be a log-factor improvement) than in the smooth case (where there is at most a constant-factor improvement, with no first-order improvement whatsoever in standard principal-agent models).

Methodologically, we develop a new technique for bounding equilibrium payoffs in repeated games with private monitoring. The starting point is that continuation payoff rewards or punishments incur an efficiency loss related to the curvature of the boundary of the feasible payoff set, while providing incentives that are proportional to a likelihood ratio difference $(p(y|a)-p(y|a'))/p(y|a)$. Since the likelihood ratio difference is a martingale increment (as the expected likelihood ratio difference under $p(\cdot|a)$ equals 0), large deviations theory can be used to bound the cumulative likelihood ratio difference over any number of periods. This bound connects the inefficiency and "incentive strength" of any strategy profile, so that any equilibrium where players do not take myopic best responses must incur a certain amount of inefficiency, regardless of whether signals are public or private.

Relation to the literature. Our finding that the gains from non-recursive equilibria are small contrasts with two strands of prior literature that find large gains. These strands share the feature that continuation value transfers are impossible with public strategies. This feature reduces the efficiency of public strategies and thereby generates large gains from non-recursive private strategies.

First, Holmström and Milgrom (1987) study a dynamic principal-agent model where the agent exerts effort over T periods and consumption occurs at the end of the game. The value of withholding feedback is large: without feedback, first-best profit can be approximated as

[2] For public games, this was already observed by Hörner and Takahashi (2016). Hörner and Takahashi observed that "It is certainly possible that regarding imperfect monitoring, allowing equilibria in private strategies could accelerate the rate of convergence beyond the results that we have derived... This is left for future research." The current paper resolves this question.

$T \to \infty$ using a review strategy that resembles the "penalty contract" of Mirrlees (1975); with feedback, optimal contracts are linear in the count of signal realizations, and profits are bounded away from the first best for all T. The key difference from our setup is that Holmström and Milgrom's model is not a repeated game (as consumption only occurs once), so efficiency cannot be improved by transferring continuation payoffs over time.[3]

Second, several papers study principal-agent problems or games that, while repeated, do not permit continuation value transfers. Abreu, Milgrom, and Pearce (1991) and Sannikov and Skrzypacz (2007) consider settings without pairwise identifiability, while Matsushima (2004) and Fuchs (2007) restrict attention to block belief-free equilibria. These settings preclude continuation value transfers, and consequently these papers all find that efficiency is attained as $\delta \to 1$ only when feedback is withheld.[4]

In past work (Sugaya and Wolitzky, 2017, 2018), we showed that the value of withholding feedback ("maintaining privacy") is large in some specific repeated and dynamic games when δ is small. For example, our 2018 paper examined how maintaining privacy can help sustain multi-market collusion. In contrast, the current paper shows that the value of privacy in repeated games is small when δ is close to 1.

We also relate to the broader literature on feedback in dynamic agency and games. We consider complete information repeated games without payoff-relevant state variables, so feedback concerns only past performance, which is payoff-irrelevant in the continuation game. In contrast, most of the literature on feedback in dynamic agency considers dynamic games with additional state variables, such as an agent's ability (Ederer, 2010; Smolin, 2021), other agents' progress in a tournament (Gershkov and Perry, 2009; Aoyagi, 2010; Ely et al., 2024), whether a project has been completed (Halac, Kartik, and Liu, 2017; Ely et al., 2023), or the evolution of an exogenous state (Ely and Szydlowski, 2020; Orlov, Skrzypacz, and Zryumov, 2020; Ball, 2023). An exception is Lizzeri, Meyer, and Persico (2002), who

[3] Relatedly, Frick, Iijima, and Ishii (2024) consider a one-shot principal-agent model and study the rate at which profit converges to the first best as the number of signal observations grows. They find that this rate is much faster for review strategies than for linear contracts.

[4] Matsushima considers two-player games where signals are conditionally independent, so each player does not learn about the status of her review. This form of lack of feedback is essential for efficiency in belief-free equilibria. Sugaya (2022) shows how mixed strategies can be used to prevent learning with conditionally dependent signals, yielding a general folk theorem under imperfect private monitoring. Rahman (2014) shows that witholding feedback restores efficiency in Sannikov and Skrzypacz's model.

examine optimal two-period agency contracts with and without a “midterm review.”

We also contribute to the literature on review strategies, introduced by Rubinstein (1979), Rubinstein and Yaari (1983), and Radner (1985), and developed by Abreu, Milgrom, and Pearce (1991) and Matsushima (2001, 2004). These papers all show that review strategies can support efficient outcomes when $\delta \to 1$ (or when there is no discounting at all). In contrast, we identify limitations of review strategies when $\delta < 1$ and show that review strategies cannot greatly outperform recursive contracts when δ is close to 1.

Methodologically, the closest papers are Hörner and Takahashi (2016), who build on Fudenberg, Levine, and Maskin's recursive methods to show that inefficiency is of order $(1-\delta)^{1/2}$ in repeated finite-action games with public monitoring; and Sugaya and Wolitzky (2023), who obtain bounds on the strength of players' equilibrium incentives in repeated finite-action games with private monitoring.[5] Rather than bounding incentives, the current paper derives a tradeoff between incentives and efficiency (e.g., program (4) below) and uses it to characterize the rate of convergence. In addition, the arguments in our 2023 paper are based on variance decomposition, while the current paper requires more precise estimates from martingale large deviations theory.

Finally, our exact characterization of first-order inefficiency in repeated principal-agent models relates to Sannikov (2008) and Sadzik and Stacchetti (2015), who derive similar results under public monitoring in continuous time or “frequent action” models. Here, our main contribution is showing that withholding feedback leaves first-order inefficiency unchanged.

Outline. The paper is organized as follows. Section 2 describes the model. Section 3 gives an informal overview of our results. Section 4 establishes general upper bounds on equilibrium efficiency. Section 5 establishes that these bounds are attainable in public equilibria (excepting a log factor in the finite-action case). Combining these results implies that the gains from non-recursive equilibria are small. Section 6 gives a stronger result for principal-agent problems. Section 7 discusses extensions.

[5] Hörner and Takahashi also consider the rate of convergence toward weakly individually rational payoff vectors, which they show can be strictly slower. We focus on strictly individually rational payoffs.

2 Preliminaries

This section introduces our model of repeated games with public monitoring and blind repeated games.

A *stage game* $G = (I, A, u)$ consists of a finite set of players $I = \{1, \ldots, N\}$, a product set of actions $A = \times_{i \in I} A_i$, and a payoff function $u_i : A \to \mathbb{R}$ for each $i \in I$. We assume that each A_i is a nonempty, compact metric space, and each u_i is continuous.[6] By the Debreu-Fan-Glicksberg theorem, the stage game admits a Nash equilibrium in mixed actions.

We fix some basic notation: the sets of stage-game Nash and correlated equilibria are $\Sigma^{NE} \subseteq \times_{i \in I} \Delta(A_i)$ and $\Sigma^{CE} \subseteq \Delta(A)$; the feasible payoff set is $F = \text{co}\left(\{u(a)\}_{a \in A}\right) \subseteq \mathbb{R}^N$; the sets of stage-game Nash and correlated equilibrium payoffs are $V^{NE} = \{v : v = u(\alpha)$ for some $\alpha \in \Sigma^{NE}\}$ and $V^{CE} = \left\{v : v = u(\alpha) \text{ for some } \alpha \in \Sigma^{CE}\right\}$; the Euclidean metric and norm on $\mathbb{R}^N$ are $d(\cdot, \cdot)$ and $\|\cdot\|$; the set of unit vectors (or *directions*) in $\mathbb{R}^N$ is $\Lambda = \{\lambda \in \mathbb{R}^N : \|\lambda\| = 1\}$; the boundary of F is bnd(F); the set of exposed points of F is exp(F); and, for any $v \in$exp(F), the set of exposing directions is $\Lambda_v = \{\lambda \in \Lambda : v = \text{argmax}_{w \in F} \lambda \cdot w\}$.[7]

A *monitoring structure* (Y, p) consists of a set of possible signal realizations Y and a family of conditional probability distributions $p(y|a)$. We assume that either Y is finite and y is drawn according to a probability mass function $p(y|a)$, or Y is a subset of a measurable space and y is drawn according to a density $p(y|a)$: we use the same notation $p(y|a)$ for both cases. We assume $p(y|a) > 0$ for all $y \in Y, a \in A$. This non-moving support assumption is crucial and, in particular, excludes perfect monitoring.

We also require that the monitoring structure satisfies the following key assumption:

Assumption 1 There exists a number $K > 0$ such that, for any $a \in A$, $i \in I$ and $a_i' \in A_i$, we have

$$\mathbb{E}^{y \sim p(\cdot|a)}\left[\exp\left(\theta \frac{p(y|a) - p(y|a_i', a_{-i})}{p(y|a)}\right)\right] \leq \exp\left(\frac{\theta^2 K}{2}\right) \qquad \text{for all } \theta \in \mathbb{R}. \tag{1}$$

[6] As is standard, we linearly extend the payoff functions u_i to distributions $\alpha \in \Delta(A)$. Here and throughout, for any compact metric space X, $\Delta(X)$ denotes the set of Borel probability measures on X, endowed with the weak* topology.

[7] Recall that, by definition, $v \in \exp(F)$ iff Λ_v is non-empty. An example at the end of Section 4.1 clarifies the necessity of considering payoff vectors that are not just extreme but exposed.

Assumption 1 says that the likelihood ratio difference between $p(\cdot|a)$ and $p(\cdot|a_i', a_{-i})$ has a *sub-Gaussian distribution*, where the number K is called a *variance proxy*.[8] For example, Assumption 1 holds if Y is finite, or if $Y \subseteq \mathbb{R}^n$ and $y = g(a) + \varepsilon$, where $g : A \to Y$ is a deterministic function and ε has a multivariate normal distribution with covariance matrix independent of a.[9] As we will see, Assumption 1 lets us apply results from large deviation theory to bound the power of tail tests that aggregate signals over many periods.

In a *repeated game with public monitoring*, in each period $t \in \mathbb{N}$, each player i takes an action a_i, and then a signal y is drawn according to $p(y|(a_i)_i)$ and is publicly observed. A history for player i at the beginning of period t takes the form $h_i^t = (a_{i,t'}, y_{t'})_{t'=1}^{t-1}$. A strategy σ_i for player i maps histories h_i^t to distributions over actions $a_{i,t}$. A strategy for player i is *public* if it depends on h_i^t only through its public component $y^t = (y_{t'})_{t'=1}^{t-1}$. Players choose strategies to maximize discounted expected payoffs, with common discount factor $\delta \in [0, 1)$. A *perfect public equilibrium* (*PPE*) is a profile of public strategies that, beginning at any period t and any public history y^t, forms a Nash equilibrium from that period on. We denote the repeated game with public monitoring with stage game G, monitoring structure (Y, p), and discount factor δ by $\Gamma^P(\delta)$, and we denote the corresponding set of PPE payoff vectors by $E^P(\delta) \subseteq \mathbb{R}^N$. Thus, $E^P(\delta)$ is the set of attainable payoffs in a (recursive) PPE where signals are publicly observed.

In a *blind repeated game*, the players are assisted by a mediator with commitment power.[10] In each period $t \in \mathbb{N}$, (i) the mediator privately recommends an action $r_i \in A_i$ to each player i, (ii) each player i takes an action a_i, and (iii) a signal y is drawn according to $p(y|(a_i)_i)$ and is observed only by the mediator. A history for the mediator at the beginning of period t takes the form $h_0^t = \left((r_{i,t'})_i, y_{t'}\right)_{t'=1}^{t-1}$, while a history for player i just before she takes an action in period t takes the form $h_i^t = \left((r_{i,t'}, a_{i,t'})_{t'=1}^{t-1}, r_{i,t}\right)$. A strategy σ_0 for the mediator maps histories h_0^t to distributions over recommendation profiles $(r_{i,t})_i$, while a strategy σ_i for player i maps histories h_i^t to distributions over actions $a_{i,t}$. We denote the blind repeated game with stage game G, monitoring structure (Y, p), and discount factor δ by $\Gamma^B(\delta)$, and

[8] See, e.g., Buldygin and Kozachenko (2000).

[9] In these cases, (1) holds with K equal to the variance of the likelihood ratio difference $(p(y|a) - p(y|a_i', a_{-i}))/p(y|a)$ (i.e., the χ^2-*divergence* of $p(|a_i', a_{-i})$ from $p(y|a)$).

[10] We previously introduced the notion of a blind repeated game in Sugaya and Wolitzky (2017, 2023).

we denote the corresponding set of Nash equilibrium payoff vectors (taking the union over all possible mediator strategies) by $E^B(\delta) \subseteq \mathbb{R}^N$. Thus, $E^B(\delta)$ is the set of attainable payoffs in a (possibly non-recursive) Nash equilibrium where signals are privately observed by a mediator.[11]

By standard arguments (similar to Forges, 1986), any Nash equilibrium outcome $\mu \in \Delta((A \times Y)^\infty)$ (i.e., any equilibrium distribution over infinite paths of action profiles and signals) in $\Gamma^P(\delta)$ can also be implemented by a Nash equilibrium in $\Gamma^B(\delta)$ where the players follow the mediator's recommendations on path. Since PPE is a refinement of Nash equilibrium, it follows that $E^P(\delta) \subseteq E^B(\delta)$. Our goal is to evaluate the advantage of arbitrary equilibria based on private signals over recursive equilibria based on public signals: that is, to assess the size of the set $E^B(\delta) \backslash E^P(\delta)$.

Remark 1 *The model is easily adapted to allow a player with commitment power (such as the principal in a standard principal-agent model) or one or more players with perfectly observed actions (such as a principal who offers contracts each period in a relational contracting model). A player with commitment power is treated like any other player, except that no incentive constraints are imposed on her strategy. For example, in a principal-agent model,* Σ^{NE} *is the set of mixed action profiles where the agent does not have a profitable deviation. Moreover, it suffices to impose non-moving support (and sub-Gaussianity) only for the agent, so that* $\operatorname{supp} p(\cdot|a) = \operatorname{supp} p(\cdot|a')$ *for all* a, a' *that agree on the principal's action. Similarly, to extend our results to the case where some players' actions are perfectly observed, let* $I^* \subseteq I$ *be the set of players with observable actions, and assume that deviations by players* $i \in I \backslash I^*$ *do not affect the support of* p*, so that* $\operatorname{supp} p(\cdot|a) = \operatorname{supp} p(\cdot|a_i', a_{-i})$ *for all* $a \in A$*,* $i \in I \backslash I^*$*, and* $a_i' \in A_i$*. Then our Theorem 1 applies for any* $v \in \exp(F)$ *that cannot be attained by an action profile distribution* α *such that* $g_i(s_i, \alpha) = 0$ *for each player* $i \in I \backslash I^*$ *and each manipulation* s_i *(where these objects are defined in Section 4.2), while our Theorem 2 applies verbatim.*

[11] Note that a player's payoff in the blind game is not measurable with respect to her own information. The blind game may thus withhold feedback from the players to an unrealistic extent—but this only strengthens our finding that withholding feedback has limited value.

3 Overview of Results

We first provide an informal overview of our results. We focus on two leading cases: finite stage games and games where the boundary of the feasible payoff set has positive curvature.

3.1 Finite Games

Our results for finite games can be illustrated in the context of a one-sided prisoner's dilemma, where the payoff matrix is

	L	R
C	$2,2$	$0,0$
D	$3,0$	$1,1$

and the monitoring structure is given by $Y = \{0,1\}^2$ and $p\left((y_1, y_2) \mid (a_1, a_2)\right) = p_1 (y_1|a_1)\, p_2 (y_2|a_2)$, where $p_1 (1|C) = p_2 (1|L) = 1/2$ and $p_1 (1|D) = p_2 (1|R) = 1/4$.[12] We investigate the possibility of attaining payoffs close to the efficient payoff vector $(2, 2)$.

First, consider PPE payoff vectors in the public game. Hörner and Takahashi (2016) showed that the minimum distance between such a vector and the efficient payoff vector $(2, 2)$ is of order $(1-\delta)^{1/2}$. This result relies on Fudenberg, Levine, and Maskin's (1994) recursive characterization of PPE and is generalized by our Theorem 2.

Next, consider arbitrary Nash equilibrium payoff vectors in the blind game. There is a wide range of non-recursive equilibria in the blind game. A leading example of these equilibria is *review strategies* (Radner, 1985; Abreu, Milgrom, and Pearce, 1991; Matsushima, 2004), which aggregate signals over T periods—during which the players take constant actions—before adjusting the players' actions. Heuristically, an optimal review strategy pools information for $T = O\left((1-\delta)^{-1}\right)$ periods and then applies a penalty if the number of "good signals" (e.g., $y_1 = 1$ for motivating cooperation in the one-sided prisoner's dilemma) over these periods falls short of a cutoff. Call the number of standard deviations by which the number of good signals falls short of its mean the *score*. Since the number of good signals, normalized by $T^{-1/2}$, is approximately normally distributed, for any cutoff score z the proba-

[12] The exact numbers are just for concreteness. We refer to them only in Proposition 1 and its proof.

bility that a single signal is pivotal is $O\left(T^{-1/2}\phi\left(z\right)\right) = O\left(\left(1-\delta\right)^{1/2}\phi\left(z\right)\right)$.[13] As stage game payoffs are $O\left(1-\delta\right)$, incentive compatibility requires that z is at most $O\left(\left(-\log\left(1-\delta\right)\right)^{1/2}\right)$, which in turn implies that the review strategy's "false positive rate" (and hence its minimum inefficiency) is $\Phi\left(-z\right) = O\left(\left(\left(1-\delta\right)/\left(-\log\left(1-\delta\right)\right)\right)^{1/2}\right)$,[14]. Thus, review strategies improve on PPE by at most a factor of $\left(-\log\left(1-\delta\right)\right)^{1/2}$. Our Theorem 1 implies that this factor is un-improvable for any finite stage game. Thus, combining Theorems 1 and 2 shows that withholding feedback accelerates convergence to efficiency by at most a factor of $\left(-\log\left(1-\delta\right)\right)^{1/2}$. Moreover, our Proposition 1 constructs an equilibrium that attains this factor in the one-sided prisoner's dilemma, which shows that this result is tight.

3.2 Positive Curvature

Now consider infinite games where the boundary of the feasible payoff set has positive curvature. In this case, Theorem 2 shows that PPE in the public game can attain inefficiency of order $1-\delta$. As we explain following the statement of Theorem 2, this reduction in inefficiency relative to finite games results because a smooth set of equilibrium payoffs can approximate a smooth set of feasible payoffs more closely than a kinked set of feasible payoffs. Conversely, Theorem 1 shows that arbitrary Nash equilibria in the blind game cannot attain inefficiency of order less than $1-\delta$. Thus, in the positive curvature case, withholding feedback does not accelerate convergence to efficiency. Moreover, our Theorem 3 shows that, in principal-agent problems, withholding feedback does not reduce first-order inefficiency.

4 Maximum Efficiency with Arbitrary Strategies

4.1 Main Result

We now turn to our formal results. Our first theorem gives an upper bound for the rate of convergence of $E^B\left(\delta\right)$ toward an exposed point $v \in \exp\left(F\right)$ that is not attainable as a static correlated equilibrium. As indicated above, the bound depends on the order of curvature of

[13] Here and throughout the paper, ϕ and Φ denote the standard normal pdf and cdf, respectively.
[14] This follows from the standard normal Mills ratio approximation: $\Phi\left(-z\right) \approx \phi\left(z\right)/z$ for $z \gg 0$.

the boundary of F at v.

Definition 1 *Fix an exposed point $v \in \exp(F)$. For any $\beta \geq 1$, the boundary of F has* max-curvature of order at least β at v *if, for all $\lambda \in \Lambda_v$, there exists $\eta > 0$ such that*

$$\lambda \cdot (v - w) \geq \eta d(v, w)^{\beta} \qquad \text{for all } w \in \text{bnd}(F).$$

The boundary of F has max-curvature of order β at v *if*

$$\beta = \inf\left\{\tilde{\beta} : \text{bnd}(F) \text{ has } \max\text{-curvature of order at least } \tilde{\beta} \text{ at } v\right\}.$$

This definition says that moving away from v in F entails an efficiency loss of order at least β, relative to Pareto weights λ. Heuristically, $\text{bnd}(F)$ is approximated by a power function of degree β at v. To understand the definition, the key cases to consider are $\beta = 1$, $\beta = 2$, and the limit case $\beta = \infty$.[15]

- The $\beta = 1$ case arises when the stage game G is finite. This implies a first-order efficiency loss from moving away from any extreme point. This case is studied by Abreu, Pearce, and Stacchetti (1990), Fudenberg, Levine, and Maskin (1994), Hörner and Takahashi (2016), and Sugaya and Wolitzky (2023).

- The $\beta = 2$ case arises when the boundary of F has positive curvature. This case is studied by Green and Porter (1984), Spear and Srivastava (1987), Sannikov (2007, 2008), and Sadzik and Stacchetti (2015). More generally, if $\beta \leq 2$ then the boundary of F has non-zero curvature: its curvature is positive but finite if $\beta = 2$ and is infinite if $\beta < 2$.

- The $\beta = \infty$ case arises when the boundary of F is linear at v. This occurs in repeated games with transferable utility, as in Athey and Bagwell (2001), Levin (2003), and Goldlücke and Kranz (2012).

[15] In addition, to appreciate the role of the "max" in the definition, suppose that $N = 2$, $(0, 0) \in F$, and the local boundary of F at $(0, 0)$ is given by $f(x) = -x$ if $x < 0$ and $f(x) = x^2$ if $x \geq 0$. Then the max-curvature of $\text{bnd}(F)$ at $(0, 0)$ is 2.

The following is our key theorem.

Theorem 1 *Fix an exposed point $v \in \exp(F) \setminus V^{CE}$ where* bnd(F) *has* max-*curvature of order β, and fix a direction $\lambda \in \Lambda_v$. Then there exists $c > 0$ such that*

$$\lambda \cdot (v - w) \geq c\zeta(\delta) \qquad \textit{for all } \delta < 1 \textit{ and } w \in E^B(\delta), \qquad \textit{where}$$
$$\zeta(\delta) = \begin{cases} \left(\frac{1-\delta}{\max\{-\log(1-\delta),1\}}\right)^{1/2} & \textit{if } \beta = 1, \\ (1-\delta)^{\max\{\beta/2,\beta-1\}} & \textit{if } \beta > 1. \end{cases}$$

The key implications of Theorem 1 are as follows:

- For Pareto weights where welfare is maximized at a kink in bnd(F), equilibrium inefficiency in the blind game is at least $O\left(\left((1-\delta)/(-\log(1-\delta))\right)^{1/2}\right)$.
- For Pareto weights where welfare is maximized at a point where bnd(F) has positive curvature, equilibrium inefficiency in the blind game is at least $O(1-\delta)$.

We will see that both of these bounds—as well as the $(1-\delta)^{\beta/2}$ bound for $\beta \in (1, 2]$—are tight. Moreover, the bound in the kinked case remains tight up to log-factor slack in the public game, while the bound in the $\beta \in (1, 2]$ case remains tight up to constant-factor slack in the public game. These results imply that the gains from non-recursive equilibria are small at any point of non-zero curvature.[16]

We outline the proof of Theorem 1 in the next subsection. The basic logic is that if a repeated game Nash equilibrium gives payoffs close to $v \in \exp(F)$, then the stage game payoff must be close to v almost all the time along the equilibrium path of play. Since signals have full support, this implies that payoffs remain close to v almost all the time even after low-probability signal realizations. This in turn implies that, on average, equilibrium continuation play does not vary much with the signal realizations. But then, if $v \notin V^{CE}$,

[16] In contrast, if welfare is maximized at a point with max-curvature of order $\beta > 2$, Theorem 1 allows inefficiency much smaller than $1-\delta$. This bound is tight in the $\beta \to \infty$ limit, since in some games with linear Pareto frontiers efficiency is exactly achieved at some $\delta < 1$ (e.g., Athey and Bagwell, 2001). We conjecture that the $(1-\delta)^{\beta-1}$ bound given by Theorem 1 is in fact tight for any $\beta > 2$—in that there exists some game and $v \in \exp(F) \setminus V^{CE}$ with max-curvature of order β that can be approached at rate $(1-\delta)^{\beta-1}$—but we have not proved this.

we can conclude that δ must be so high that even small variations in continuation play can provide strong incentives.[17]

We mention a couple technical aspects of the statement of Theorem 1. First, generically, the condition $v \in \exp(F) \setminus V^{CE}$ is equivalent to $v \in \exp(F) \setminus V^{NE}$: since v is extremal, the distinction only matters in the non-generic case where v is attained at two different pure action profiles. Second, the condition $\lambda \in \Lambda_v$ (i.e., $v = \text{argmax}_{w \in F} \lambda \cdot w$) cannot be weakened to $v \in \text{argmax}_{w \in F} \lambda \cdot w$. To see this, consider the stage game

	L	R
C	$1, 1$	$0, 1$
$D1$	$2, 0$	$-2, 0$
$D2$	$-2, 0$	$2, 0$

Here the point $v = (1, 1)$ is exposed and is not attainable as a static CE; $\text{bnd}(F)$ has curvature of order 1 (i.e., a kink) at v; and $v \in \text{argmax}_{w \in F} \lambda \cdot w$ for $\lambda = (0, 1)$. But the point $w = (0.5, 1)$ is attained by the static NE $\left(C, \frac{1}{2}L + \frac{1}{2}R\right)$ (so $w \in E^B(\delta)$ for all $\delta \in [0, 1)$) and satisfies $\lambda \cdot w = \lambda \cdot v$, so the conclusion of Theorem 1 fails.

4.2 Proof Sketch for Theorem 1

We sketch the proof of Theorem 1, deferring the details to the appendix. Fix any $v \in \exp(F) \setminus V^{CE}$ and $\lambda \in \Lambda_v$. We wish to derive a lower bound for $\lambda \cdot (v - w)$—the inefficiency of w in direction λ—which holds for any $w \in E^B(\delta)$.

We introduce some notation. Note that any outcome $\mu \in \Delta((A \times Y)^\infty)$ defines a marginal distribution over period-t action profiles, $\alpha_t^\mu \in \Delta(A)$, as well as an *occupation measure* $\alpha^\mu \in \Delta(A)$, defined as

$$\alpha^\mu = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} \alpha_t^\mu.$$

[17] This logic is the same as that of Theorem 6.5 of Fudenberg, Levine, and Maskin (1994) (who credit Madrigale, 1986), which says that an extremal non-static Nash payoff vector v cannot be exactly attained for any $\delta < 1$ under full-support monitoring. Fudenberg, Levine, and Maskin state this result for PPE, but the same argument works for Nash. Theorem 1 is a quantitative version of this result.

The players' ex ante payoffs under μ are determined by α^{μ}, as, by linearity of u,

$$(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}u\left(\alpha_t^{\mu}\right)=u\left((1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\alpha_t^{\mu}\right)=u\left(\alpha^{\mu}\right).$$

Thus, letting $\mathcal{M}^B(\delta)$ be the set of Nash equilibrium outcomes μ in the blind game $\Gamma^B(\delta)$, we wish to derive a lower bound for

$$\inf_{\mu\in\mathcal{M}^B(\delta)}\lambda\cdot\left(v-u\left(\alpha^{\mu}\right)\right).$$

Now, for each player i, let S_i denote the set of functions $s_i : A_i \to A_i$, which we call *manipulations*. For any $i \in I$, $\alpha \in \Delta(A)$, and $s_i \in S_i$, define the *deviation gain*

$$g_i\left(s_i,\alpha\right)=\sum_{a\in A}\alpha(a)\left(u_i\left(s_i\left(a_i\right),a_{-i}\right)-u_i(a)\right).$$

The interpretation is: if the recommended action profile a is drawn according to α and player i takes $s_i(a_i)$ when recommended a_i rather than obeying the recommendation, her expected payoff gain is $g_i(s_i,\alpha)$. Finally, for any complete history of play $h=(a_t,y_t)_{t=1}^{\infty}$ and any player i and manipulation s_i, let

$$\hat{u}_{i,t}(h)=u_i\left(a_t\right)-v_i \qquad \text{and} \qquad \ell_{i,t}\left(s_i,h\right)=\frac{p\left(y_t|a_t\right)-p\left(y_t|s_i\left(a_{i,t}\right),a_{-i,t}\right)}{p\left(y_t|a_t\right)}.$$

That is, $\hat{u}_t(h)$ is the difference between player i's realized period t payoff at history h and v_i—which we will call player i's period t *reward* at history h—and $\ell_t(h)$ is the realized likelihood ratio difference of the period t signal y_t at the period t action profile a_t, as compared to the action profile $(s_i(a_{i,t}),a_{-i,t})$ that results when player i manipulates according to s_i.

A simple necessary condition for an outcome μ to be consistent with equilibrium play (Lemma 6 in the appendix) is that, for each player i, manipulation s_i, and period t, we have

$$g_i\left(s_i,\alpha_t^{\mu}\right)\leq\mathbb{E}^{\mu}\left[\ell_{i,t}\left(s_i,h\right)\sum_{t'=t+1}^{\infty}\delta^{t'-t}\hat{u}_{i,t'}(h)\right]. \tag{2}$$

This inequality holds because, if it were violated, player i could gain by obeying her recommendation in every period other than t, while manipulating according to s_i in period t. Given this inequality, since $\mathrm{bnd}(F)$ has max-curvature of order β at v, we have

$$\begin{aligned} & \inf_{\mu \in \mathcal{M}^B(\delta)} \lambda \cdot (v - u(\alpha^\mu)) \\ \geq \quad & \inf_{\mu \in \Delta((A \times Y)^\infty)} \sup_{i \in I, s_i \in S_i} \inf_{\substack{(\hat{u}_{i,t}(h))_{t,h} \in [-\bar{u}_i, \bar{u}_i] \\ \text{s.t. } (2)}} \mathbb{E}^\mu \left[(1-\delta) \sum_{t=1}^\infty \delta^{t-1} \eta \left| \hat{u}_{i,t}(h) \right|^\beta \right], \end{aligned} \tag{3}$$

where $\bar{u}_i$ is the range of u_i. Intuitively, the program (3) minimizes the maximum over players i and manipulations s_i of the β^{th} moment of the deviation of player i's stage game payoff from v_i, subject to the incentive constraint (2).

To prove the theorem, it remains to bound (3) as a function of δ and β. To do so, consider the inner problem where μ is fixed and $(i, s_i) \in \mathrm{argmax}_{i,s_i} g_i(s_i, \alpha^\mu)$. Let $(1-\delta)\delta^{t-1}\xi_t$ denote the Lagrange multiplier on the period t incentive constraint, and form the Lagrangian

$$\sup_{(\xi_t)_t \geq 0} \inf_{(\hat{u}_t(h))_{t,h} \in [-\bar{u}, \bar{u}]} (1-\delta) \sum_{t=1}^\infty \delta^{t-1} \mathbb{E}^\mu \left[\eta \left| \hat{u}_t(h) \right|^\beta + \xi_t \left(g_t^\mu - \ell_t(h) \sum_{t'=t+1}^\infty \delta^{t'-t} \hat{u}_{t'}(h) \right) \right], \tag{4}$$

where we have simplified notation by letting $\hat{u}_t(h) = \hat{u}_{i,t}(h)$, $g_t^\mu = g_i(s_i, \alpha_t^\mu)$, and $\ell_t(h) = \ell_{i,t}(s_i, h)$. The Lagrangian expresses a tradeoff between efficiency and incentives: to maximize efficiency, the reward $\hat{u}_t(h)$ must minimize the sum of the inefficiency resulting from the curvature of $\mathrm{bnd}(F)$ (i.e., $\eta \left| \hat{u}_t(h) \right|^\beta$) and an incentive cost in each earlier period $\tilde{t} < t$ (i.e., $-\xi_{\tilde{t}} \ell_{\tilde{t}}(h) \hat{u}_t(h)$). Moreover, if we take ξ_t to be constant across periods (i.e., $\xi_t = \xi \; \forall t$), we can reverse the order of summation between t and t' (and also note that $\hat{u}_1(h) = 0$ for all h at the optimum) to rewrite the Lagrangian as

$$\sup_{\xi \geq 0} \inf_{(\hat{u}_t(h))_{t \geq 2, h} \in [-\bar{u}, \bar{u}]} (1-\delta) \sum_{t=2}^\infty \delta^{t-1} \mathbb{E}^\mu \left[\eta \left| \hat{u}_t(h) \right|^\beta - \xi \mathcal{L}_{t-1}(h) \hat{u}_t(h) \right] + \xi g^\mu,$$

where $\mathcal{L}_t(h) = \sum_{t'=1}^t \ell_{t'}(h)$ and $g^\mu = g_i(s_i, \alpha^\mu) = (1-\delta) \sum_{t=1}^\infty \delta^{t-1} g_t^\mu$. Thus, to bound the Lagrangian, it suffices to bound the probability that $|\mathcal{L}_t(h)|$ is large. Since $\mathcal{L}_t$ is a martingale with sub-Gaussian increments (by Assumption 1), the required bounds follow from large deviations theory (Lemma 9). Intuitively, these bounds say that sequences of signals with large cumulative likelihood ratio differences—which are highly informative when they occur—also occur with low equilibrium probability, and hence do not provide a large amount of information on average. These bounds imply that the value of the Lagrangian—and hence inefficiency under outcome μ—cannot be too much smaller than g^{μ}. Finally, since $v \in \exp(F) \setminus V^{CE}$, if $u(\alpha^{\mu})$ is close to v then $g^{\mu} = \max_{i,s_i} g_i(s_i, \alpha^{\mu})$ is not too small (Lemma 4), which yields the desired bound.

4.3 Tightness of the Efficiency Bound in the Kinked Case

We will see in the next section that inefficiency of order $(1-\delta)^{\beta/2}$ is attainable when $\beta \in [1,2]$ under public monitoring. This implies that the lower bound on inefficiency in Theorem 1 cannot be improved when $\beta \in (1,2]$ (the smooth, non-zero curvature case). Here we show that, in the kinked case ($\beta = 1$), inefficiency of order $((1-\delta)/-\log(1-\delta))^{1/2}$ is sometimes attainable in the blind game. This shows that the lower bound on inefficiency in Theorem 1 also cannot be improved when $\beta = 1$. Consequently, withholding feedback can accelerate the rate of convergence by at most a factor of $(-\log(1-\delta))^{-1/2}$ in the kinked case.

Our result here concerns the one-sided prisoner's dilemma described in Section 3.1

Proposition 1 *In the one-sided prisoner's dilemma, there exists $c > 0$ such that, for any sufficiently large $\delta < 1$, there exists $v \in E^B(\delta)$ satisfying*

$$v_1 = v_2 > 2 - c\left(\frac{1-\delta}{-\log(1-\delta)}\right)^{1/2}.$$

The proof constructs a review strategy with inefficiency of order $((1-\delta)/(-\log(1-\delta)))^{1/2}$, as sketched in Section 3.1.

5 Attainable Efficiency with Public Strategies

We now show that the maximum efficiency levels identified in Theorem 1 are attainable under public monitoring in the smooth, non-zero curvature case, and are attainable up to a log factor in the kinked case. To this end, denote the set of feasible and strictly individually rational payoffs by $F^* = \left\{v \in F : v_i > \underline{v}_i := \min_{\alpha_{-i} \in \times_{j \neq i} \Delta(A_j)} \max_{a_i \in A_i} u_i(a_i, \alpha_{-i}) \;\; \forall i\right\}$. For $v \in \text{bnd}(F^*)$, define $\Lambda_v^* = \{\lambda \in \Lambda : v \in \text{argmax}_{w \in F^*} \lambda \cdot w\}$.

The following definition is a counterpart of Definition 1, adjusted to apply to all boundary points rather than only exposed points. It says that moving away from v along the boundary of F^* entails an efficiency loss of order at most β, relative to Pareto weights λ.[18]

Definition 2 *Fix a boundary point $v \in \text{bnd}(F^*)$. For any $\beta \geq 1$, the boundary of F^* has* min-curvature of order at most β *at v if, for all $\lambda \in \Lambda_v^*$, there exists $k > 0$ such that*

$$\lambda \cdot (v - w) < k d(v, w)^{\beta} \qquad \textit{for all } w \in \text{bnd}(F^*).$$

The boundary of F^* has min-curvature of order β at v *if*

$$\beta = \sup\left\{\tilde{\beta} : \text{bnd}(F^*) \;\; \textit{has} \;\; \min\textit{-curvature of order at most } \tilde{\beta} \textit{ at } v\right\}.$$

Note that, at any exposed point $v \in \text{bnd}(F^*)$, the min-curvature of $\text{bnd}(F^*)$ is at least 1 and at most the max-curvature.

The following assumption generalizes standard identification conditions for the public-monitoring folk theorem to the case where action sets can be infinite.

Assumption 2 There exists $\bar{x} > 0$ such that the following conditions hold:

i. For each i, there exists a minmax profile against i, $\alpha^i \in \times_{j \neq i} \Delta(A_j) \times A_i$, and $x_j : Y \to [-\bar{x}, \bar{x}]$ for each $j \neq i$, such that

$$a_j \in \underset{a_j'}{\text{argmax}}\, u_j\left(a_j', \alpha_{-j}\right) + \mathbb{E}\left[x_j(y) \,|a_j', \alpha_{-j}\right] \;\text{ for all } j \neq i \text{ and } a_j \in \text{supp}(\alpha_j). \quad (5)$$

ii. For each $a \in A$, $c \in \{-1, +1\}$, and (i, j) with $i \neq j$, there exists $x_i : Y \to [-\bar{x}, \bar{x}]$ such

[18] For example, if $N = 2$, $(0, 0) \in F^*$, and the local boundary of F^* at $(0, 0)$ is given by $f(x) = -x$ if $x < 0$ and $f(x) = x^2$ if $x \geq 0$, then the min-curvature of $\text{bnd}(F^*)$ at $(0, 0)$ is 1.

that

$$a_i \in \underset{a_i'}{\operatorname{argmax}}\, u_i\left(a_i', a_{-i}\right) + \mathbb{E}\left[x_i\left(y\right)|a_i', a_{-i}\right] \qquad \text{and} \tag{6}$$

$$a_j \in \underset{a_j'}{\operatorname{argmax}}\, \mathbb{E}\left[c x_i\left(y\right)|a_j', a_{-j}\right]. \tag{7}$$

Intuitively, Assumption 2 says that, when payoff transfers of magnitude at most $\bar{x}$ are available, players $-i$ can be incentivized to minmax player i, and player i can be incentivized to take a given action a_i via transfers from player j without affecting player j's incentive to take a given action a_j.[19]

We consider the rate of convergence of $E^P(\delta)$ toward a strictly individually rational payoff vector $v \in \operatorname{bnd}(F^*)$. For finite stage games, Hörner and Takahashi (2016) show that this rate equals $(1-\delta)^{1/2}$. Thus, withholding feedback can accelerate the rate of convergence by at most a factor of $(-\log(1-\delta))^{-1/2}$ in finite-action games. We now show that whenever the boundary of F^* has non-zero curvature ($\beta \leq 2$), the rate equals $(1-\delta)^{\beta/2}$. (We discuss the zero curvature case below.) Thus, withholding feedback cannot accelerate the rate of convergence in smooth games with non-zero curvature.

We require the standard assumption that $\dim F^* = N$ and further exclude payoff vectors where some player obtains her maximum feasible payoff.[20]

Theorem 2 *Assume that Assumption 2 holds and* $\dim F^* = N$*, and fix any* $v \in \operatorname{bnd}(F^*)$*, satisfying* $v_i < \max_a u_i(a)$ *for all* i*, where* $\operatorname{bnd}(F^*)$ *has* min-*curvature of order* $\beta \geq 1$*. Then there exists* $c > 0$ *such that* $d\left(v, E^P(\delta)\right) \leq c(1-\delta)^{\min\{\beta,2\}/2}$ *for any sufficiently large* $\delta < 1$*.*

Theorem 2 builds on Fudenberg, Levine and Maskin (1994), Hörner and Takahashi (2016), and Sugaya and Wolitzky (2023). As these authors showed, a given level of inefficiency

[19] Assumption 2 is similar to assumptions (A1)–(A3) of Kandori and Matsushima (1998). The difference is that we allow $|A| = \infty$ and state Assumption 1 directly in terms of the existence of transfers x that satisfy (5)–(7), while Kandori and Matsushima assume that $|A| < \infty$ and hence can state conditions in terms of the convex hull of the set of vectors of signal probabilities generated by different actions, which imply the existence of transfers x satisfying (5)–(7) by the separating hyperplane theorem.

[20] To attain the same rate of convergence toward "max points," one must show that, for $v \in \operatorname{argmax}_{v' \in F^*} \lambda \cdot v'$, as λ approaches a coordinate direction e_i, v must be implemented by action profiles where player i's deviation gain vanishes. In finite-action games, Hörner and Takahashi (2016) show that this is possible under a genericity condition on payoffs. For a class of infinite-action games (the linear-concave games considered below), this is possible under a bounded cross-partial derivative condition.

relative to an exposed point v and a direction $\lambda \in \Lambda_v$ is attainable under public monitoring if it equals the distance in direction λ between v and a self-generating ball $B \subseteq F$. We must thus find a self-generating ball $B \subseteq F$ at distance $O\left((1-\delta)^{\min\{\beta,2\}/2}\right)$ to v in direction $\lambda \in \Lambda_v$. To this end, let $d = d(B, v)$ be the desired distance, and (without loss) let $u = v - d\lambda$ be the closest point to v in B. (See Figure 1.) Consider decomposing u into an instantaneous payoff v and continuation payoffs $(w(y))_y$ that lie on the translated tangent hyperplane H with normal vector λ passing through the point $\mathbb{E}[w(y)] = v - ((1-\delta)/\delta) d\lambda$. Under Assumption 2, the continuation payoffs $(w(y))_y$ can be chosen to enforce v on $H \cap B$ if the diameter of $H \cap B$, which we denote by x, is of order $1-\delta$. At the same time, denoting the radius of B by r, the Pythagorean theorem gives $(x/2)^2 + (r - ((1-\delta)/\delta) d)^2 = r^2$, and hence $x = O\left(\sqrt{(1-\delta) rd}\right)$. It follows that the product rd is of order $1-\delta$, and hence $r = O\left((1-\delta)^{1-\min\{\beta,2\}/2}\right)$. Finally, for a point v where the (max-)curvature of $\operatorname{bnd}(F)$ equals β, a ball B with radius $r = O\left((1-\delta)^{1-\min\{\beta,2\}/2}\right)$ and center $v - (r+d)\lambda$, where $d = O\left((1-\delta)^{\min\{\beta,2\}/2}\right)$, lies entirely within F. For example, if $\beta = 1$ then r and d are both $O\left((1-\delta)^{1/2}\right)$ and thus shrink at the same rate as $\delta \to 1$; while if $\beta \geq 2$ then $r = O(1)$ and $d = O(1-\delta)$, so B simply shifts toward v as $\delta \to 1$.[21]

In light of Theorem 1, when $\beta > 2$ one might hope to find conditions under which $d\left(v, E^P(\delta)\right) = O\left((1-\delta)^{\beta-1}\right)$. While this may be possible, we do not pursue such a result here. The difficulty is that the corresponding ball B would have to have radius r of at least $O\left((1-\delta)^{2-\beta}\right)$ (as rd must be at least $O(1-\delta)$). While such a ball can satisfy the self-generation condition $B \subseteq F$ in a neighborhood of v, its radius explodes as $\delta \to 1$ (when $\beta > 2$), so it must violate self-generation at some point far from v. Therefore, any conditions that ensure that $d\left(v, E^P(\delta)\right)$ is less than $O(1-\delta)$ must involve the global geometry of the feasible payoff set. Investigating such conditions is left for future work.

We finally mention a class of infinite games where Assumption 2(ii) holds.[22] Say that the game is *linear-concave* if (i) for each i, A_i is a compact interval $\left[\underline{A}_i, \bar{A}_i\right] \subseteq \mathbb{R}$, and $u_i(a_i, a_{-i})$ is differentiable and concave in a_i for every a_{-i} with a bounded derivative: there exists $\kappa > 0$

[21] For the detailed argument for any β, see Lemma 17 in the online appendix.

[22] If Assumption 2(ii) holds, then even if Assumption 2(i) fails a Nash-threat folk theorem still holds, i.e., Theorem 2 holds with F^* replaced by the set of feasible payoffs that Pareto dominate a convex combination of static Nash payoffs.

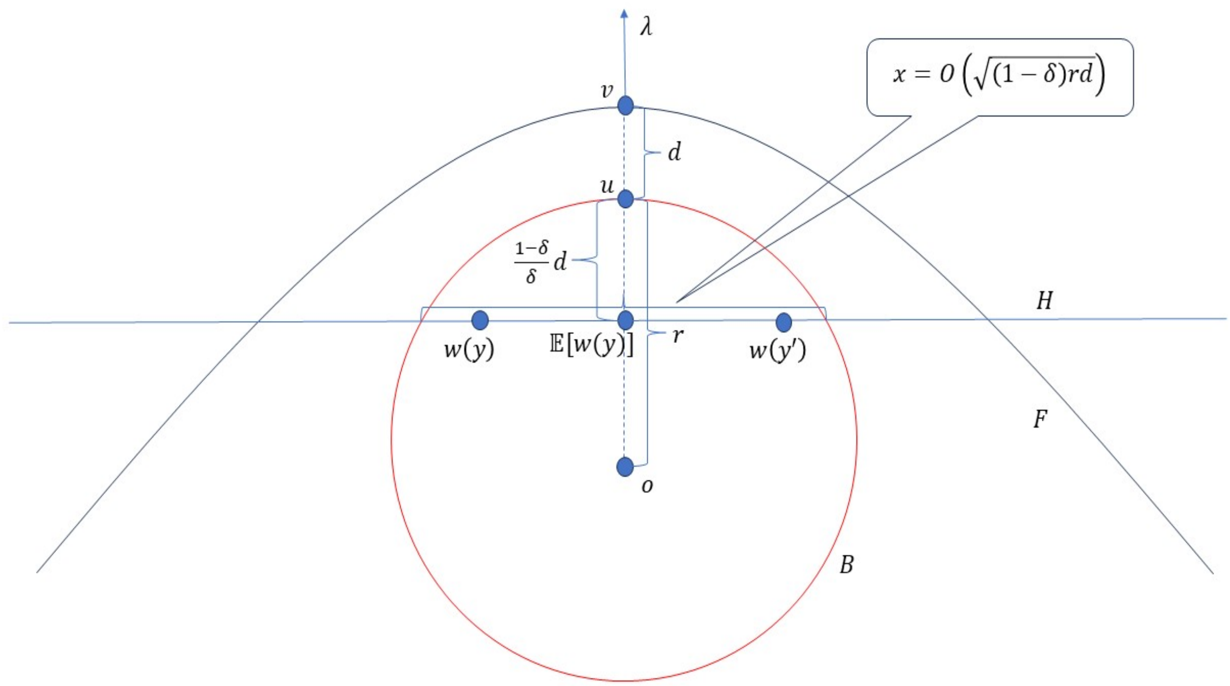

Figure 1: Self-Generating a Ball. To maximize efficiency, r and d must be chosen to minimize d subject to the constraints that $B \subseteq F$ and x is at least $O(1-\delta)$.

such that $|\partial u_i(a_i, a_{-i})/\partial a_i| \leq \kappa$ for all i, a; and (ii) the public signal is a D-dimensional real variable, $Y = \times_{d=1}^{D} Y^d \subseteq \mathbb{R}^D$, and $\mu^d(a) = \mathbb{E}\left[y^d|a\right]$ is a linear function of a for each dimension d. In a linear-concave game, let $M^i(a) = \left(\frac{d}{da_i}\mu^d(\hat{a})\Big|_{\hat{a}=a}\right)_d$ be a D-dimensional vector representing the sensitivity of the mean public signal to player i's action. Say that a linear-concave game satisfies *pairwise identifiability* if for any a and $i \neq j$, $M^i(a) \neq 0$ and the spans of $M^i(a)$ and $M^j(a)$ intersect only at the origin.[23]

Proposition 2 *In any linear-concave game satisfying pairwise identifiability, Assumption 2(ii) holds.*

5.1 Proof Sketch for Theorem 2

We recall a key definition and lemma from Abreu, Pearce, and Stacchetti (1990).

Definition 3 *A bounded set $W \subseteq \mathbb{R}^N$ is* self-generating *if for all $\hat{v} \in W$, there exist $\alpha \in \times_i \Delta(A_i)$ and $w: Y \to \mathbb{R}^N$ satisfying*

[23] This condition is the same as Assumption 1 of Sannikov (2007).

Promise keeping (PK) $\hat{v} = (1-\delta) u(\alpha) + \delta \int_y w(y) p(y|\alpha) dy$.

Incentive compatibility (IC) $\text{supp}(\alpha_i) \subseteq \text{argmax}_{a_i} (1-\delta) u_i(a_i, \alpha_{-i}) + \delta \int_y w_i(y) p(y|a_i, \alpha_{-i}) dy$ *for all* i.

Self-generation (SG) $w(y) \in W$ *for all* y.

When (PK), (IC), and (SG) hold, we say that the pair (α, w) decomposes $\hat{v}$ on W.

Lemma 1 *Any bounded, self-generating set* W *is contained in* $E^P(\delta)$.

It thus suffices to find a bounded, self-generating set W such that $d(v, W) = O\left((1-\delta)^{\beta^*/2}\right)$, where $\beta^* = \min\{\beta, 2\}$. To do so, we first establish a sufficient condition for a ball B to be self-generating. This condition builds on Fudenberg and Levine (1994) and Sugaya and Wolitzky (2023).[24]

Definition 4 *The* maximum score in direction $\lambda \in \Lambda$ with reward bound $\bar{x} > 0$ *is*

$$k(\lambda, \bar{x}) := \sup_{\alpha \in \times_i \Delta(A_i), x: Y \to \mathbb{R}^N} \lambda \cdot \left(u(\alpha) + \int_y x(y) p(y|\alpha) dy \right), \qquad \textit{subject to}$$

1. (IC): $\text{supp}(\alpha_i) \subseteq \text{argmax}_{a_i} u_i(a_i, \alpha_{-i}) + \int_y x_i(y) p(y|a_i, \alpha_{-i}) dy$ *for all* i.

2. Half-space decomposability with reward bound $\bar{x}$ (HS$\bar{x}$): $\lambda \cdot x(y) \leq 0$ *and* $\|x(y)\| \leq \bar{x}$ *for all* y.

Lemma 2 *For any* $\bar{x} > \max_{u,u' \in F} \|u - u'\|$ *and* $\varepsilon > 0$, *if a ball* B *of radius* r *satisfies*

$$k(\lambda, \bar{x}) \geq \max_{v' \in B} \lambda \cdot v' + \varepsilon \quad \textit{for all } \lambda \in \Lambda, \quad \textit{and} \tag{8}$$

$$\bar{x}^2 \leq \frac{\delta}{1-\delta} \frac{\varepsilon r}{36}, \tag{9}$$

then B *is self-generating.*

[24] Lemma 2 is similar to Lemma 6 of Sugaya and Wolitzky (2023), but is simpler because the monitoring structure varies together with δ in our 2023 paper but is fixed in the current paper, so here we do not need as much control over the relationship between δ and the reward bound $\bar{x}$.

We then show that there exists B with $d(v,B) = O\left((1-\delta)^{\beta^*/2}\right)$ that satisfies the sufficient condition for self-generation just given.

Lemma 3 *There exist* $\bar{x} > \max_{u,u' \in F} \|u - u'\|$, $c > 0$ *and* $\bar{\delta} < 1$ *such that, for any* $\delta > \bar{\delta}$, *there exist* $\varepsilon > 0$ *and a ball* B *of radius* r *satisfying (8), (9), and* $d(v,B) \le c(1-\delta)^{\beta^*/2}$.

The proof of Lemma 3 uses Assumption 2 and the assumptions that $\dim F^* = N$, $v_i \in (\underline{v}_i, \max_a u_i(a))$ for all i, and $\mathrm{bnd}(F^*)$ has min-curvature of order $\beta \ge 1$ at v. The logic is similar to that accompanying Figure 1.

The proofs of Lemmas 2 and 3 are deferred to the online appendix. Given these lemmas, taking $\bar{x}$, c, and $\bar{\delta}$ as in Lemma 3 establishes Theorem 2.

6 A Stronger Result for the Principal-Agent Problem

In this section, we establish that withholding feedback in a standard repeated principal-agent problem leaves unchanged not only the rate of convergence to efficiency (the order of inefficiency in $1-\delta$), but also the exact level of first-order inefficiency (the constant multiplying $1-\delta$). This stronger result also has the virtue of identifying the precise features of the stage game and the monitoring structure that determine the level of first-order inefficiency.

Consider a standard repeated principal-agent problem in discrete time. In each period t, an agent chooses an effort level a from a compact interval $A = [0, \bar{A}]$, and a signal y is then drawn according to a pmf or pdf $p(y|a)$. Assume that $p(y|a)$ is twice continuously differentiable in a, with first and second derivatives $p_a(y|a)$ and $p_{aa}(y|a)$. A *contract* specifies, for each period t, a recommended effort level $r_t \in A$ as a function of the history of past recommendations and signals $(r_{t'}, y_{t'})_{t'=1}^{t-1}$, as well as the agent's current consumption $c_t \ge 0$ as a function of $(r_{t'}, y_{t'})_{t'=1}^{t}$. In the public game, the agent chooses her period t action a_t as a function of $\left((r_{t'}, a_{t'}, y_{t'})_{t=1}^{t-1}, r_t\right)$; in the blind game, she chooses a_t as a function of $\left((r_{t'}, a_{t'})_{t=1}^{t-1}, r_t\right)$ only. The agent's payoff in period t is $u(c_t) - \psi(a_t)$, where the consumption utility u is twice continuously differentiable on $\mathbb{R}_+$ with $u(0) = 0$, $u' > 0$, $u'' < 0$, $\lim_{c \to \infty} u'(c) = 0$, and

$$\sup_{c \in [0,\infty)} \frac{u''(c)}{(u'(c))^3} < 0, \tag{10}$$

and the effort cost ψ is twice continuously differentiable on A with $\psi(0) = \psi'(0) = 0$ and $\psi'' > 0$. (We discuss the role of condition (10) below.) The principal's payoff in period t is $a_t - c_t$.[25] The parties have the same discount factor $\delta \in [0, 1)$.

For any effort level $a \in A$, the *score* of the signal y is

$$\nu(y|a) = \frac{p_a(y|a)}{p(y|a)} \qquad \text{for all } y \in Y,$$

and the *Fisher information*—the variance of the score—is

$$\mathcal{I}(a) = \int_y \frac{p_a(y|a)^2}{p(y|a)} dy.$$

We require the following technical assumption, which implies that the Fisher information is finite, strictly positive, and Lipschitz continuous in a; the distribution of the score is sub-Gaussian; and a second-order condition holds.

Assumption 3 The following hold:

i. For all $a \in A$, there exists $\Delta > 0$ such that

$$\int_y \frac{\max_{\tilde{a} \in [a, a+\Delta]} p_a(y|\tilde{a})^2}{p(y|a)} dy < \infty.$$

ii. $\mathcal{I}(a)$ is strictly positive and Lipschitz continuous on A.

iii. The score $\nu(y|a)$ is sub-Gaussian with variance proxy $\mathcal{I}(a)$:

$$\int_y \exp(\theta \nu(y|a)) p(y|a) dy \leq \exp\left(\frac{\theta^2 \mathcal{I}(a)}{2}\right) \qquad \text{for all } \theta \in \mathbb{R}.$$

[25] As indicated above, a_t is not contractable. The interpretation is that $a_t - c_t$ is the principal's expected payoff, where her realized payoff is determined by y_t and c_t.

iv. There exists $\hat{K}$ such that, for all $a, \hat{a} \in A$, we have

$$\int_y \nu(y|a)\, p_{aa}(y|\hat{a})\, dy \leq 0 \qquad \text{and} \tag{11}$$

$$\int_y \frac{p_{aa}(y|\hat{a})^2}{p(y|a)} dy \leq \hat{K}. \tag{12}$$

For example, Assumption 3 is satisfied if Y is finite, or if $Y \subseteq \mathbb{R}^n$ and $y = g(a) + \varepsilon$ for a deterministic function $g : A \to Y$ with a bounded gradient and multivariate normal noise ε with covariance independent of a. Note that Assumption 3(iii) strengthens Assumption 1.

For any $w \in \left[-\psi\left(\bar{A}\right), \bar{u}\right)$, where $\bar{u} = \lim_{c\to\infty} u(c) \in \mathbb{R}_+ \cup \{\infty\}$, let $\bar{F}(w)$ be the first-best payoff for the principal when the agent's payoff equals w, which is given by

$$\bar{F}(w) = \max_{a \in A} a - u^{-1}(w + \psi(a)).$$

Let $\bar{a}(w)$ be the maximizer (which is unique, as the maximand is strictly concave), and let $\bar{c}(w) = u^{-1}(w + \psi(\bar{a}(w)))$ be the corresponding consumption for the agent. Note that $\bar{F}$ is twice continuously differentiable, and $\bar{a}$ and $\bar{c}$ are continuously differentiable. In addition, since $\psi'(0) = 0$ and $u' > 0$, we have $\bar{a}(w) > 0$ for all $w \in \left[-\psi\left(\bar{A}\right), \bar{u}\right)$.

Finally, let $F_\delta^B(w)$ (resp., $F_\delta^P(w)$) denote the maximum payoff for the principal over all $v \in E^B(\delta)$ (resp., $v \in E^P(\delta)$) where the agent's payoff is w. That is, $F_\delta^B(w)$ is the principal's second-best payoff in the blind game, while $F_\delta^P(w)$ is her second-best payoff in the public game. Recall that $E^B(\delta) \supseteq E^P(\delta)$, so $F_\delta^B(w) \geq F_\delta^P(w)$. Nonetheless, we show that $F_\delta^B(w)$ and $F_\delta^P(w)$ agree up to a first-order approximation as $\delta \to 1$.

Theorem 3 *For any $\delta < 1$ and $w \in (0, \bar{u})$, we have*

$$F_\delta^B(w) = \bar{F}(w) + \frac{1-\delta}{\delta} \frac{\psi'(\bar{a}(w))^2}{\mathcal{I}(\bar{a}(w))} \frac{\bar{F}''(w)}{2} + o(1-\delta) \qquad \textit{and}$$

$$F_\delta^P(w) = \bar{F}(w) + \frac{1-\delta}{\delta} \frac{\psi'(\bar{a}(w))^2}{\mathcal{I}(\bar{a}(w))} \frac{\bar{F}''(w)}{2} + o(1-\delta), \tag{13}$$

where in each equation $o(1-\delta)$ stands for a (different) function satisfying $\lim_{\delta\to 1} o(1-\delta)/(1-\delta) = 0$.

Theorem 3 shows that, whether or not the agent receives feedback, the first-order inefficiency of an optimal contract is precisely

$$\frac{1-\delta}{\delta}\frac{\psi'\left(\bar{a}\left(w\right)\right)^{2}}{\mathcal{I}\left(\bar{a}\left(w\right)\right)}\frac{\bar{F}''\left(w\right)}{2}. \tag{14}$$

In the public monitoring case, this result is similar to Theorem 5 of Sannikov (2008) and Corollary 1 of Sadzik and Stacchetti (2015), but we consider a discrete-time game with a general monitoring structure, while Sannikov considers a continuous-time game with Brownian noise, and Sadzik and Stacchetti consider the "frequent action limit" of a class of monitoring structures that converge to Brownian noise. However, the key point of Theorem 3 is that the first-order inefficiency is exactly the same under private monitoring. Thus, even if the principal can conceal the agent's past performance in a standard repeated principal-agent problem, she can do little better than to fully reveal it and utilize a public contract.

A rough intuition for Theorem 3 is that, with high probability, the agent's continuation payoff is approximately constant for a long time under an optimal contract, so there is little information about the continuation payoff to conceal, and thus little value from concealing it.

The proof of Theorem 3 is facilitated by the principal's ability to commit to delivering any feasible promised continuation value for the agent. It may be possible to generalize Theorem 3 to smooth games with 1-dimensional actions and product structure monitoring (as considered by Sannikov, 2007), but this would require constructing equilibria that attain specific continuation payoff vectors far from the initial target vector. This possibility is left for future research.

We finally comment on the role of condition (10). This condition implies that the second-order efficiency loss from varying the agent's utility is uniformly bounded away from zero. With CRRA utility $u\left(c\right) = c^{1-\gamma}/\left(1-\gamma\right)$, it holds iff $\gamma \geq 1/2$. Without this condition, review strategies with infrequent, large rewards may yield a first-order improvement over (14) if $u'\left(c\right)$ converges to 0 sufficiently slowly as $c \to \infty$.

6.1 Proof Sketch for Theorem 3

We first bound $F_{\delta}^{B}(w)$ from above. Fix a period t and a small constant $\Delta > 0$, and consider the manipulation where, whenever the agent is recommended effort a in period t, she instead takes effort $a - (\psi'(a)/\mathcal{I}(a))\Delta$. (The agent thus shades her effort more after recommendations where effort is more costly or less detectable.) For this manipulation to be unprofitable for all t and $\Delta > 0$, we must have

$$\mathbb{E}^{\mu}\left[\frac{\psi'(a_t)^2}{\mathcal{I}(a_t)} - \frac{\psi'(a_t)\,\nu(y_t|a_t)}{\mathcal{I}(a_t)}\sum_{t'=t+1}^{\infty}\delta^{t'-t}\hat{u}_{t'}\right] \leq 0 \qquad \text{for all } t, \tag{15}$$

where μ is the equilibrium occupation measure and $\hat{u}_{t'} = u(c_{t'}) - \psi(a_{t'}) - w$ is the deviation of the agent's period t' utility from w (see Lemma 22 in the online appendix). Letting $(1-\delta)\delta^{t-1}\xi_t$ denote the Lagrange multiplier on this relaxed period t incentive constraint (as in the proof Theorem 1) and letting $\xi_t = \frac{1-\delta}{\delta}\bar{F}''(w)$ for all t, the inner problem in the Lagrangian (4) becomes

$$\inf_{(\hat{u}_t(h))_{t,h}} (1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\mathbb{E}^{\mu}\left[\begin{array}{c}\bar{F}(w+\hat{u}_t(h)) - \bar{F}(w) \\ +\frac{1-\delta}{\delta}\bar{F}''(w)\left(\frac{\psi'(a_t)^2}{\mathcal{I}(a_t)} - \frac{\psi'(a_t)\nu(y_t|a_t)}{\mathcal{I}(a_t)}\sum_{t'=t+1}^{\infty}\delta^{t'-t}\hat{u}_{t'}(h)\right)\end{array}\right].$$

Taking the Taylor approximation $\bar{F}(w+\hat{u}_t(h)) - \bar{F}(w) \approx \frac{\bar{F}''(w)}{2}\hat{u}_t(h)^2$ and rearranging, this equals

$$\frac{1-\delta}{\delta}\frac{\bar{F}''(w)}{2}\inf_{(\hat{u}_t(h))_{t,h}}\sum_{t=1}^{\infty}\mathbb{E}^{\mu}\left[\begin{array}{c}\delta^t\hat{u}_t(h)^2 - 2\frac{1-\delta}{\delta}\frac{\psi'(a_t)\nu(y_t|a_t)}{\mathcal{I}(a_t)}\sum_{t'=t+1}^{\infty}\delta^{t'}\hat{u}_{t'}(h) \\ +2(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\frac{\psi'(a_t)^2}{\mathcal{I}(a_t)}\end{array}\right].$$

The FOC for $\hat{u}_t$ is

$$\hat{u}_t(h) = \frac{1-\delta}{\delta}\sum_{t'=1}^{t-1}\frac{\psi'(a_{t'})\,\nu(y_{t'}|a_{t'})}{\mathcal{I}(a_{t'})} \qquad \text{for all } t \geq 2 \text{ and } h.^{26}$$

[26] A subtlety is that the Taylor approximation for $\bar{F}(w+\hat{u}_t) - \bar{F}(w)$ is slack if $\hat{u}_t$ is large, which occurs if some score $\nu(y_{t'}|a_{t'})$ is large. However, large scores occur with low probability by Assumption 3, and the efficiency gain from relying on large scores is limited by condition (10). These complications are addressed in the proof of Lemma 10 in the appendix.

Substituting the FOC into the Lagrangian (together with $\hat{u}_1(h) = 0$ for all h) and simplifying (see equation (30)) gives

$$\frac{1-\delta}{\delta}\mathbb{E}^{\alpha^\mu}\left[\frac{\psi'(a)^2}{\mathcal{I}(a)}\right]\frac{\bar{F}''(w)}{2}.$$

Finally, to attain inefficiency of order $1-\delta$, we must have $|\alpha^\mu - a(\bar{w})| \leq O(1-\delta)$. Inefficiency is thus no less than

$$\begin{aligned}
&\frac{1-\delta}{\delta}\frac{\psi'(\bar{a}(w))^2}{\mathcal{I}(\bar{a}(w))}\frac{\bar{F}''(w)}{2} + \frac{1-\delta}{\delta}\left(\mathbb{E}^{\alpha^\mu}\left[\frac{\psi'(a)^2}{\mathcal{I}(a)}\right] - \frac{\psi'(\bar{a}(w))^2}{\mathcal{I}(\bar{a}(w))}\right)\frac{\bar{F}''(w)}{2} \\
= &\frac{1-\delta}{\delta}\frac{\psi'(\bar{a}(w))^2}{\mathcal{I}(\bar{a}(w))}\frac{\bar{F}''(w)}{2} + o(1-\delta).
\end{aligned}$$

We next bound $F_\delta^P(w)$ from below. Given a continuation payoff w_t for the agent, suppose the principal implements first-best effort $\bar{a}(w_t)$ by offering the corresponding first-best consumption $\bar{c}(w_t)$ and providing incentives entirely by varying the continuation payoff w_{t+1} while making it a martingale: $\mathbb{E}[w_{t+1}|w_t] = w_t$. The Taylor approximation of inefficiency is then equal to

$$\sum_{t=1}^{\infty}\mathbb{E}^\mu\left[(w_{t+1}-w_t)^2\right]\frac{\bar{F}''(w_t)}{2}.$$

To bound the variance $\mathbb{E}^\mu\left[(w_{t+1}-w_t)^2\right]$, note that the agent's incentive constraint is

$$\bar{a}(w_t) \in \underset{a}{\operatorname{argmax}}\,(1-\delta)(u(\bar{c}(w_t)) - \psi(a)) + \delta\int p(y|a)(w_{t+1}(y) - w_t)\,dy,$$

with FOC

$$(1-\delta)\psi'(\bar{a}(w_t)) + \delta\int p_a(y|\bar{a}(w_t))(w_{t+1}(y) - w_t)\,dy = 0.$$

To minimize variance subject to the agent's FOC, the principal takes $w_{t+1}(y) - w_t$ proportional to $\nu(y|\bar{a}(w_t))$, which gives variance

$$\mathbb{E}^\mu\left[(w_{t+1}-w_t)^2\right] = \left(\frac{1-\delta}{\delta}\right)^2\frac{\psi'(\bar{a}(w_t))^2}{\mathcal{I}(\bar{a}(w_t))}.^{27}$$

[27] It is infeasible to take $w_{t+1}(y) - w_t$ exactly proportional to $\nu(y|\bar{a}(w_t))$ when $\nu(y|\bar{a}(w_t))$ is large, but this is a rare event by Assumption 3. See Lemma 25 in the online appendix for the formal construction of the agent's continuation payoff.

The resulting ex ante inefficiency thus equals

$$
\begin{aligned}
& \mathbb{E}^{\mu}\left[\sum_{t}\delta^{t-1}\left(\frac{1-\delta}{\delta}\right)^{2}\frac{\psi'\left(\bar{a}\left(w_{t}\right)\right)^{2}}{I\left(\bar{a}\left(w_{t}\right)\right)}\frac{\bar{F}''\left(w_{t}\right)}{2}\right] \\
= \; & \frac{1-\delta}{\delta^{2}}\frac{\psi'\left(\bar{a}\left(w\right)\right)^{2}}{\mathcal{I}\left(\bar{a}\left(w\right)\right)}\frac{\bar{F}''\left(w\right)}{2}+\sum_{t}\delta^{t-1}\left(\frac{1-\delta}{\delta}\right)^{2}\mathbb{E}^{\mu}\left[\frac{\psi'\left(\bar{a}\left(w_{t}\right)\right)^{2}}{I\left(\bar{a}\left(w_{t}\right)\right)}\frac{\bar{F}''\left(w_{t}\right)}{2}-\frac{\psi'\left(\bar{a}\left(w\right)\right)^{2}}{\mathcal{I}\left(\bar{a}\left(w\right)\right)}\frac{\bar{F}''\left(w\right)}{2}\right] \\
= \; & \frac{1-\delta}{\delta}\frac{\psi'\left(\bar{a}\left(w\right)\right)^{2}}{\mathcal{I}\left(\bar{a}\left(w\right)\right)}\frac{\bar{F}''\left(w\right)}{2}+o\left(1-\delta\right),
\end{aligned}
$$

where the last line follows as $\mathbb{E}^{\mu}\left[\left(w_{t+1}-w_{t}\right)^{2}\right]=O\left(\left(1-\delta\right)^{2}\right)$, and hence $\left|w_{t}-w\right|\leq O\left(\sqrt{t}\left(1-\delta\right)\right)$ with high probability, so in the second line both the sum from $t=1$ to $\left(1-\delta\right)^{-1}$ and the sum from $t=\left(1-\delta\right)^{-1}$ to ∞ are $o\left(1-\delta\right)$ (see the proof of Lemma 12).

7 Discussion

7.1 The Low-Discounting/Low-Monitoring Double Limit

This paper focuses on the rate at which inefficiency vanishes as $\delta\to1$ for a fixed monitoring structure. In contrast, in Sugaya and Wolitzky (2023) we showed that in the double limit where $\delta\to1$ at the same time as monitoring precision degrades, the prospects for cooperation depend on a ratio of discounting and monitoring precision. This double limit arises, for example, in the "frequent action limit" considered by Abreu, Milgrom and Pearce (1991), Fudenberg and Levine (2007), Sannikov and Skrzypacz (2010), and Sadzik and Stacchetti (2015), where signals are parameterized by an underlying continuous-time process, actions and signal observations occur simultaneously every Δ units of time, and the analysis concerns the $\Delta\to0$ limit.

The results of the current paper extend to the low-discounting/low-monitoring double limit. To see this, maintain the assumption that the monitoring structure is sub-Gaussian with variance proxy K, but now view K as a variable that varies simultaneously with the discount factor. Since K proxies the variance of the likelihood ratio difference, a higher value for K corresponds to more precise monitoring, so the low-discounting/low-monitoring double limit arises when $K\to0$ and $\delta\to1$ simultaneously. For example, in the standard frequent action limit, discounting and monitoring vanish at the same rate, so $(1-\delta)/K$ remains constant as $K \to 0$ and $\delta \to 1$.

From this more general perspective, it can be shown (by nearly the same proof) that Theorem 1 holds verbatim with $(1-\delta)/K$ in place of $1-\delta$. Conversely, Theorem 2 also holds with $(1-\delta)/K$ in place of $1-\delta$, under the condition that $\bar{x}$ in Assumption 2 can be taken to be of order $K^{-1/2}$. For example, this condition holds with finite signals with $p(y|a)$ bounded away from zero, or with Gaussian signals.

7.2 Summary and Directions for Future Research

This paper has used a rate-of-convergence approach to analyze the gains from non-recursive equilibria in standard repeated agency problems and games with patient players. The main result is that these gains are "small": (i) in finite-action games, non-recursive equilibria reduce inefficiency by at most a log factor; (ii) in smooth games, non-recursive equilibria reduce inefficiency by at most a constant factor; and (iii) in smooth principal-agent problems, non-recursive equilibria do not reduce first-order inefficiency. The key force underlying these results is that, while pooling information across periods leads to more precise monitoring, it also entails larger rewards and punishments, which reduces the scope for providing incentives by transferring continuation value rather than destroying it.

A basic lesson of our analysis is that the value of withholding feedback in dynamic agency is very different in a one-off production process that unfolds gradually over time (as in Holmström and Milgrom, 1987) as compared to a repeated interaction. Since continuation payoff transfers are impossible in one-shot interactions, the monitoring benefit of withholding feedback dominates, so withholding feedback can be very valuable. But in repeated interactions, this benefit is offset by the cost of using larger rewards and punishments, which limit continuation payoff transfers.

We mention some possible extensions of our results. First, as discussed in Section 5, characterizing rates of convergence toward exposed points with curvature of order $\beta > 2$ is a challenging open question involving non-local properties of the feasible payoff set. Second, as discussed in Section 6, it may be possible to generalize Theorem 3 from smooth agency problems to smooth games. Third, it would be interesting to relax the assumption that the likelihood ratio difference is sub-Gaussian. This could result in a faster rate of convergence, because rare but highly informative signals would become more common, and such signals become more useful as δ increases. Fourth, it remains to characterize the rate of convergence when different actions of player 1 generate signals of player 2's action of very unequal precision (e.g., in games where player 2's action is entirely unobserved unless player 1 pays a "monitoring cost"). Fifth, it would be interesting to extend our results to stochastic games. More broadly, the rate of convergence to efficiency as discounting vanishes many be a useful lens for analyzing a range of other questions about long-run economic relationships, beyond the value of withholding performance feedback.

Appendix

A Proof of Theorem 1

We first bound a player's deviation gain at any $\alpha \in \Delta(A)$ that attains payoffs close to v.

Lemma 4 *There exist $\varepsilon > 0$ and $\gamma > 0$ such that, for all $\alpha \in \Delta(A)$ satisfying $\lambda \cdot (v - u(\alpha)) < \varepsilon$, there exist a player i and a manipulation s_i such that $g_i(\alpha, s_i) > \gamma$.*

Proof. Since $v \in \exp(F) \backslash V^{CE}$, for all $\alpha \in \Delta(A)$ such that $v = u(\alpha)$, there exist i and s_i such that $g_i(s_i, \alpha) > 0$. Let

$$\gamma = \frac{1}{2} \inf_{\alpha \in \Delta(A): v = u(\alpha)} \sup_{i, s_i} g_i(s_i, \alpha).$$

Note that $\gamma > 0$. To see this, note that $g_i(\mathrm{Id}, \alpha) = 0$ for all i, α, so $\gamma \geq 0$, and suppose toward a contradiction that there exists a sequence α^n such that $v = u(\alpha^n)$ for all n and $\sup_{i, s_i} g_i(s_i, \alpha^n) \to 0$. Since $\Delta(A)$ is weak*-compact by Alaoglu's theorem, taking a subsequence if necessary, $\alpha^n \to \alpha \in \Delta(A)$. Moreover, since each u_i is continuous, $u(\alpha) = v$; and since each A_i is compact, by the maximum theorem, $\sup_{s_i} g_i(s_i, \alpha) = \lim_n \sup_{s_i} g_i(s_i, \alpha^n) = 0$ for all i, contradicting $v \notin V^{CE}$.

Now suppose that for all $\varepsilon > 0$ there exists $\alpha^\varepsilon \in \Delta(A)$ satisfying $\lambda \cdot (v - u(\alpha^\varepsilon)) < \varepsilon$ and $g_i(s_i, \alpha^\varepsilon) < \gamma$ for all i, s_i. Taking a subsequence if necessary, $\alpha^\varepsilon \to \alpha \in \Delta(A)$. Moreover, we have $u(\alpha) = \lim_\varepsilon u(\alpha^\varepsilon) = v$ (since $u(\alpha^\varepsilon) \in F$ and $v \in \exp(F)$), and $\sup_{s_i} g_i(s_i, \alpha) = \lim_\varepsilon \sup_{s_i} g_i(s_i, \alpha^\varepsilon) \leq \gamma$ for all i (by the maximum theorem), so $\sup_{i,s_i} g_i(s_i, \alpha) \leq \gamma$, contradicting the definition of γ. ■

Fix such ε and γ. Next, for any outcome μ and period T, define the occupation measure over the first T periods by $\alpha^{\mu,T} = \left((1-\delta)/\left(1-\delta^T\right)\right) \sum_{t=1}^T \delta^{t-1} \alpha_t^\mu$, and define $T(\delta) = \lceil (\log 2)/(-\log \delta) \rceil$. We first bound $\lambda \cdot (v - u(\alpha^\mu))$ for any μ where all player's deviation gains over the first $T(\delta)$ periods are small.

Lemma 5 *For any outcome μ where $g_i\left(s_i, \alpha^{\mu,T(\delta)}\right) \leq \gamma$ for all players i and manipulations s_i, we have $\lambda \cdot (v - u(\alpha^\mu)) \geq \varepsilon/2$.*

Proof. Since $\delta^T \leq 1/2$ by construction, we have

$$\begin{aligned} \lambda \cdot (v - u(\alpha^\mu)) &= (1-\delta) \sum_{t=1}^{\infty} \delta^{t-1} \lambda \cdot (v - u(\alpha_t^\mu)) \\ &\geq (1-\delta) \sum_{t=1}^{T} \delta^{t-1} \lambda \cdot (v - u(\alpha_t^\mu)) = \left(1-\delta^T\right) \lambda \cdot \left(v - u\left(\alpha^T\right)\right) \geq \frac{\lambda}{2} \cdot \left(v - u\left(\alpha^T\right)\right). \end{aligned}$$

By construction of (ε, γ), if $\lambda \cdot \left(v - u\left(\alpha^T\right)\right) < \varepsilon$ then $\sup_{i,s_i} g_i\left(\alpha^T, s_i\right) > \gamma$. Hence, $\sup_{i,s_i} g_i\left(\alpha^T, s_i\right) \leq \gamma$ implies $\lambda \cdot \left(v - u\left(\alpha^T\right)\right) \geq \varepsilon$, as desired. ■

We next establish the incentive constraint, (2).

Lemma 6 *For any equilibrium outcome $\mu \in \mathcal{M}^B(\delta)$, player i, manipulation s_i, and period t, we have $g_i(s_i, \alpha_t^\mu) \leq \mathbb{E}^\mu\left[\ell_{i,t}(s_i, h) \sum_{t'=t+1}^\infty \delta^{t'-t} \hat{u}_{t'}(h)\right]$.*

Proof. For any sequence of action profiles $(a_t)_{t=1}^\infty$ and any period t, let $w_t(h) = \sum_{t'=t}^\infty \delta^{t'-t} u_i(a_{t'})$. Since μ is an equilibrium outcome, for every $t \in \mathbb{N}$ we have

$$g_i(\alpha_t^\mu, s_i) \leq \int_{h^t, a_t, y_t} (p(y_t|a_t) - p(y_t|s_i(a_{i,t}), a_{-i,t})) \delta \mathbb{E}\left[w_{t+1}(h) | h^t, a_t, y_t\right] d\mu\left(h^t, a_t\right) dy_t.$$

This holds because, if she follows her recommendation in every period $t' \neq t$ while manipulating according to s_i in period t, player i obtains an expected continuation payoff of $\int_{h^t, a_t, y_t} p(y_t|s_i(a_{i,t}), a_{-i,t}) \mathbb{E}[w_{t+1}(h) | h^t, a_t, y_t] d\mu(h^t, a_t) dy_t$ in period $t+1$, and this deviation must be unprofitable. The lemma follows as

$$\begin{aligned}
& \int_{h^t,a_t,y_t} \left(p\left(y_t|a_t\right) - p\left(y_t|s_i\left(a_{i,t}\right),a_{-i,t}\right)\right) \delta\mathbb{E}\left[w_{t+1}\left(h\right)|h^t,a_t,y_t\right] d\mu\left(h^t,a_t\right) dy_t \\
= & \int_{h^t,a_t,y_t} p\left(y_t|a_t\right) \ell_{i,t}\left(s_i,h\right) \delta\mathbb{E}\left[w_{t+1}\left(h\right)|h^t,a_t,y_t\right] d\mu\left(h^t,a_t\right) dy_t \\
= & \int_h \ell_{i,t}\left(s_i,h\right) \sum_{t'=t+1}^{\infty} \delta^{t'-t}\hat{u}_{t'}\left(h\right) d\mu\left(h\right),
\end{aligned}$$

where the last line follows by iterated expectation. ■

We now come to our key lemma, which bounds (4)—and hence $\lambda \cdot (v - u(\alpha^\mu))$—for any μ where some player's deviation gain over the first $T(\delta)$ periods is large.

Lemma 7 *There exists $\bar{c} > 0$ such that, for any outcome μ, player i, and manipulation s_i, and discount factor $\delta < 1$ satisfying $g_i\left(s_i, \alpha^{\mu,T(\delta)}\right) > \gamma$, we have*

$$\sup_{(\xi_t)_t \geq 0} \inf_{(\hat{u}_t(h))_{t,h} \in [-\bar{u},\bar{u}]} (1-\delta) \sum_{t=1}^{\infty} \delta^{t-1} \mathbb{E}^\mu \left[\eta \left|\hat{u}_t\left(h\right)\right|^\beta + \xi_t \left(g_t^\mu - \ell_t\left(h\right) \sum_{t'=t+1}^{\infty} \delta^{t'-t}\hat{u}_{t'}\left(h\right)\right)\right] \geq \bar{c}\zeta\left(\delta\right).$$

Together, Lemmas 5, 6, and 7 imply that $\lambda \cdot (v - u(\alpha^\mu)) \geq \max\{\varepsilon/2, \bar{c}\zeta(\delta)\} \geq \max\{\varepsilon/2, \bar{c}\}\zeta(\delta)$ for all $\delta < 1$ and $\mu \in \mathcal{M}^B(\delta)$. Theorem 1 therefore holds with $c = \min\{\varepsilon/2, \bar{c}\}$.

It thus remains to prove Lemma 7. To this end, let $\xi_t = \xi \geq 0$ if $t \leq T(\delta)$, and $\xi_t = 0$ otherwise. Letting $T = T(\delta)$ to ease notation, we then have

$$\begin{aligned}
(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\xi_t\ell_t\left(h\right)\sum_{t'=t+1}^{\infty}\delta^{t'-t}\hat{u}_{t'}\left(h\right) &= (1-\delta)\,\xi\sum_{t=1}^{T}\ell_t\left(h\right)\sum_{t'=t+1}^{\infty}\delta^{t'-1}\hat{u}_{t'}\left(h\right) \\
&= (1-\delta)\,\xi\sum_{t=2}^{\infty}\delta^{t-1}\mathcal{L}_{\min\{t-1,T\}}\hat{u}_t\left(h\right), \qquad \text{and} \\
(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\xi_t g_t^\mu &= \xi\,(1-\delta)\sum_{t=1}^{T}\delta^{t-1}g_t^\mu = \xi\left(1-\delta^T\right)g_i\left(s_i,\alpha^{\mu,T}\right) \geq \frac{\xi\gamma}{2}.
\end{aligned}$$

In total, we see that (4) is no less than

$$\sup_{\xi \geq 0}\left(\frac{\xi\gamma}{2} + \inf_{(\hat{u}_t(h))_{t\geq 2,h} \in [-\bar{u},\bar{u}]} (1-\delta)\sum_{t=2}^{\infty}\delta^{t-1}\mathbb{E}^\mu\left[\eta\left|\hat{u}_t\left(h\right)\right|^\beta - \xi\mathcal{L}_{\min\{t-1,T\}}\hat{u}_t\left(h\right)\right]\right). \tag{16}$$

The following lemma thus establishes Lemma 7.

Lemma 8 *For each $\beta \geq 1$, there exists $\bar{c} > 0$ such that, for any μ and δ, the value of (16) is no less than $\bar{c}\zeta(\delta)$.*

In turn, Lemma 8 relies on the following large deviations bound for martingales.

Lemma 9 *Let $(X_t)_{t\geq 1}$ be a sequence of martingale increments adapted to a filtration $(H_t)_{t\geq 0}$, so that $\mathbb{E}[X_t|H_{t-1}] = 0$, and let $(\omega_t)_{t\geq 1}$ be a stochastic process adapted to the same filtration satisfying $\mathbb{E}[\exp(\theta X_t)|H_{t-1}] \leq \exp(\theta^2\omega_t/2)$ for all $t \geq 1$ and $\theta \in \mathbb{R}$. Let $S_T = \sum_{t=1}^T X_t$ and $W_T = \sum_{t=1}^T \omega_t$. For all $T \geq 1$, we have $\mathbb{E}[\exp(\theta S_T)] \leq \exp(\theta^2 W_T/2)$, and hence (i) $\Pr(|S_T| \geq x) \leq 2\exp(-x^2/(2W_T))$ for all $x \geq 0$, and (ii) $\mathbb{E}[|S_T|^\varphi] \leq 2(\varphi W_T/e)^{\varphi/2}$ for all $\varphi \geq 0$.*

Proof. By iterated expectation,

$$\mathbb{E}[\exp(\theta S_T)] = \mathbb{E}[\exp(\theta S_{T-1})\mathbb{E}[\exp(\theta X_T)|H_{T-1}]] \leq \mathbb{E}[\exp(\theta S_{T-1})]\exp\left(\theta^2\omega_T/2\right).$$

Recursively applying the same argument gives $\mathbb{E}[\exp(\theta S_T)] \leq \exp\left(\theta^2 W_T/2\right)$. Applying the Chernoff bound then gives (i) and (ii): see, e.g., Lemmas 1.3 and 1.4 of Buldygin and Kozachenko (2000). ■

Proof of Lemma 8. We consider separately the cases where $\beta = 1$ and $\beta > 1$.

Case 1: When $\beta = 1$, the minimand in (16) is linear in $\hat{u}_t(h)$. Minimizing over $\hat{u}_t(h) \in [-\bar{u}, \bar{u}]$, we see that (16) equals

$$\sup_{\xi\geq 0}\left(\frac{\xi\gamma}{2} + (1-\delta)\sum_{t=2}^{\infty}\delta^{t-1}\int_{h:\xi\left|\Sigma_{t'=1}^{\min\{t-1,T\}}\ell_{t'}(h)\right|\geq\eta}\left(\eta - \xi\left|\mathcal{L}_{\min\{t-1,T\}}\right|\right)\bar{u}d\mu(h)\right). \tag{17}$$

Note that $\mathbb{E}\left[\ell_t|(a_{t'}, y_{t'})_{t'=1}^{t-1}, a_t\right] = 0$ and $\mathbb{E}\left[\exp(\theta\ell_t)|(a_{t'}, y_{t'})_{t'=1}^{t-1}, a_t\right] \leq \exp\left(\theta^2 K/2\right)$ for all t, h, and θ, by (1). Hence, by Lemma 9, $\mathcal{L}_{\min\{t-1,T\}}$ is sub-Gaussian with variance proxy $K\min\{t-1,T\}$, and thus satisfies

$$\Pr\left(\xi\left|\mathcal{L}_{\min\{t-1,T\}}\right| \geq x\right) \leq 2\left(\exp\left(-\frac{x^2}{2\xi^2 K\min\{t-1,T\}}\right)\right).$$

We thus have

$$
\begin{aligned}
& \int_{h:\xi \mathcal{L}_{\min\{t-1,T\}} \geq \eta} \left(\eta - \xi \left|\mathcal{L}_{\min\{t-1,T\}}\right|\right) d\mu\left(h\right) \\
= \;& \Pr{}^{\mu}\left(\xi \left|\mathcal{L}_{\min\{t-1,T\}}\right| \geq \eta\right)\eta - \mathbb{E}^{\mu}\left[\mathbf{1}\left\{\xi \left|\mathcal{L}_{\min\{t-1,T\}}\right| \geq \eta\right\} \left|\xi \mathcal{L}_{\min\{t-1,T\}}\right|\right] \\
= \;& -\int_{x \geq \eta} \Pr{}^{\mu}\left(\xi \left|\mathcal{L}_{\min\{t-1,T\}}\right| \geq x\right) dx \geq -2\int_{x\geq\eta} \exp\left(-\frac{x^2}{2K\xi^2 \min\left\{t-1,T\right\}}\right) dx,
\end{aligned}
$$

where the second equality is by integration by parts. Now note that

$$
\begin{aligned}
\int_{x\geq\eta} \exp\left(-\frac{x^2}{2K\xi^2\min\left\{t-1,T\right\}}\right)dx &= \sqrt{2\xi^2 K}\sqrt{\min\left\{t-1,T\right\}} \int_{y \geq \frac{\eta}{\sqrt{2\xi^2 K}\sqrt{\min\{t-1,T\}}}} \exp\left(-y^2\right) dy \\
&\leq \frac{2\xi^2 K}{\eta} \min\left\{t-1,T\right\} \exp\left(-\frac{\eta^2}{2\xi^2 K \min\left\{t-1,T\right\}}\right) \\
&\leq \frac{2\xi^2 KT}{\eta} \exp\left(-\frac{\eta^2}{2\xi^2 KT}\right), \qquad (18)
\end{aligned}
$$

where the first inequality uses the Mills ratio inequality $\phi\left(-x\right)/\Phi\left(-x\right) \geq x$ for $x \geq 0$. Hence, (17) is no less than

$$
\sup_{\xi\geq 0}\left(\frac{\xi\gamma}{2} - \frac{4\bar{u}\xi^2 K}{\eta}\left(1-\delta\right)\sum_{t=2}^{\infty}\delta^{t-1}T\exp\left(-\frac{\eta^2}{2\xi^2 KT}\right)\right) \geq \sup_{\xi\geq 0}\xi\left(\frac{\gamma}{2} - \frac{4\bar{u}\xi KT}{\eta}\exp\left(-\frac{\eta^2}{2\xi^2 KT}\right)\right).
$$

Finally, letting

$$
\xi^* = \eta\left(KT\max\left\{\log\left(2^8 KT\bar{u}^2\gamma^{-2}\right),1\right\}\right)^{-1/2},
$$

we have

$$
\begin{aligned}
& \sup_{\xi\geq 0}\xi\left(\frac{\gamma}{2} - \frac{4\bar{u}\xi KT}{\eta}\exp\left(-\frac{\eta^2}{2\xi^2 KT}\right)\right) \\
\geq \;& \xi^*\left(\frac{\gamma}{2} - 4\bar{u}\sqrt{KT}\left(\max\left\{\log\left(2^8 KT\bar{u}^2\gamma^{-2}\right),1\right\}\right)^{-1/2}\exp\left(-\frac{1}{2}\max\left\{\log\left(2^8 KT\bar{u}^2\gamma^{-2}\right),1\right\}\right)\right) \\
\geq \;& \xi^*\left(\frac{\gamma}{2} - 4\bar{u}\sqrt{KT}\exp\left(-\frac{1}{2}\log\left(2^8 KT\bar{u}^2\gamma^{-2}\right)\right)\right) \\
= \;& \frac{\xi^*\gamma}{4} \\
= \;& \frac{\eta\gamma}{4}\left(KT\max\left\{\log\left(2^8 KT\bar{u}^2\gamma^{-2}\right),1\right\}\right)^{-1/2} \\
\geq \;& \frac{\eta\gamma}{4}\left(\left(4K\log 2\right)\max\left\{\log\left(2^9\left(\log 2\right)K\bar{u}^2\gamma^{-2}\right),1\right\}\right)^{-1/2}\zeta\left(\delta\right),
\end{aligned}
$$

where the last inequality follows because $T \leq \lceil (\log 2) / (1-\delta) \rceil \leq 2 (\log 2) / (1-\delta)$ and $\max\{-\log xy, 1\} \leq 2 \max\{-\log x, 1\} \max\{-\log y, 1\}$ for all $x, y \in \mathbb{R}$. This is a constant multiple of $\zeta(\delta)$, as desired.

Case 2: When $\beta > 1$, the minimand in (16) is convex in $\hat{u}_t(h)$. Relaxing the constraint $\hat{u}_t(h) \in [-\bar{u}, \bar{u}]$ and minimizing over $\hat{u}_t(h) \in \mathbb{R}$ gives

$$\hat{u}_t(h) = \left(\frac{\xi}{\eta\beta}\right)^{\frac{1}{\beta-1}} \operatorname{sign}\left(\mathcal{L}_{\min\{t-1,T\}}\right) \left|\mathcal{L}_{\min\{t-1,T\}}\right|^{\frac{1}{\beta-1}} \qquad \text{for all } t \geq 2.$$

Hence, substituting for $\hat{u}_t(h)$, (16) is no less than

$$\sup_{\xi \geq 0} \left(\frac{\xi\gamma}{2} - \xi^{\frac{\beta}{\beta-1}} \left(\frac{1}{\eta} \frac{\beta^{\beta-1} - 1}{\beta^{\beta}} \right)^{\frac{1}{\beta-1}} (1-\delta) \sum_{t=2}^{\infty} \delta^{t-1} \mathbb{E}\left[\left|\mathcal{L}_{\min\{t-1,T\}}\right|^{\frac{\beta}{\beta-1}} \right] \right). \tag{19}$$

By Lemma 9, $\mathcal{L}_{\min\{t-1,T\}}$ is sub-Gaussian with variance proxy $K \min\{t-1, T\} \leq K(t-1)$, and thus satisfies

$$\begin{aligned}
\mathbb{E}^{\mu}\left[\left|\mathcal{L}_{\min\{t-1,T\}}\right|^{\frac{\beta}{\beta-1}} \right] &\leq 2 \left(\frac{\beta}{e(\beta-1)} \right)^{\frac{\beta}{2(\beta-1)}} (K(t-1))^{\frac{\beta}{2(\beta-1)}} \\
&\leq 2 \left(\frac{\beta}{e(\beta-1)} \right)^{\frac{\beta}{2(\beta-1)}} K^{\frac{\beta}{2(\beta-1)}} (t-1)^{\max\left\{\frac{\beta}{2(\beta-1)}, 1\right\}}.
\end{aligned}$$

Next, for any $\vartheta \geq 1$, we let $k(\vartheta) \geq 1$ satisfy

$$\sum_{t=1}^{\infty} \delta^t t^{\vartheta} \leq \frac{k(\vartheta)}{(1-\delta)^{\vartheta+1}} \qquad \text{for all } \delta. \tag{20}$$

(The existence of such $k(\vartheta)$ follows from the standard fact that $\sum_{t=1}^{\infty} \delta^t t^{\vartheta} = \Gamma(\vartheta+1)(1-\delta)^{-(\vartheta+1)} + O\left((1-\delta)^{-\vartheta}\right)$: see, e.g., Wood, 1992, eqn. (6.4).) With this definition, we have

$$\begin{aligned}
&(1-\delta) \sum_{t=2}^{\infty} \delta^{t-1} \mathbb{E}\left[\left|\mathcal{L}_{\min\{t-1,T\}}\right|^{\frac{\beta}{\beta-1}} \right] \\
\leq\ & 2 \left(\frac{\beta}{e(\beta-1)} \right)^{\frac{\beta}{2(\beta-1)}} K^{\frac{\beta}{2(\beta-1)}} k\left(\max\left\{ \tfrac{\beta}{2(\beta-1)}, 1 \right\} \right) (1-\delta)^{-\max\left\{\frac{\beta}{2(\beta-1)}, 1\right\}} \\
\leq\ & 2 \left(\frac{\beta}{e(\beta-1)} \right)^{\frac{\beta}{2(\beta-1)}} K^{\frac{\beta}{2(\beta-1)}} k\left(\max\left\{ \tfrac{\beta}{2(\beta-1)}, 1 \right\} \right) (1-\delta)^{-\max\left\{\frac{\beta}{2(\beta-1)}, 1\right\}}.
\end{aligned}$$

Thus, (19) is no less than

$$\sup_{\xi\geq 0}\xi\left(\frac{\gamma}{2}-2\left(\frac{\beta}{e\left(\beta-1\right)}\right)^{\frac{\beta}{2(\beta-1)}}\left(\frac{\beta^{\beta-1}-1}{\eta\beta^{\beta}}\right)^{\frac{1}{\beta-1}}K^{\frac{\beta}{2(\beta-1)}}k\left(\max\left\{\frac{\beta}{2(\beta-1)},1\right\}\right)\left(\frac{\xi}{(1-\delta)^{\max\{\beta/2,\beta-1\}}}\right)^{\frac{1}{\beta-1}}\right).$$

Since the coefficient of $\left(\xi/\left(1-\delta\right)^{\max\{\beta/2,(\beta-1)\}}\right)^{\frac{1}{\beta-1}}$ is independent of δ, there exists $\hat{c}>0$ such that if $\xi=4\hat{c}\left(1-\delta\right)^{\max\{\beta/2,\beta-1\}}$ then the resulting value is no less than $\xi\gamma/4$, which is again a constant multiple of $\zeta\left(\delta\right)$. ■

B Proof of Proposition 1

Consider a review strategy where the game is divided into blocks of T consecutive periods. Let $T=\lfloor\rho/(1-\delta)\rfloor$, where $\rho>0$ is a small number to be determined: note that $\rho\approx 1-\delta^{T}$ when $\delta\approx 1$. In the first block, the players are prescribed (C,L) in every period. At the end of the first block—as well any subsequent block where (C,L) is prescribed—the mediator records the summary statistic

$$E=\mathbf{1}\left\{\frac{1}{\sqrt{T}}\sum_{t=1}^{T}\left(2y_{1,t}-1\right)\leq-\sqrt{-\log\left(1-\delta\right)}\right\}.$$

(Here periods are numbered from the start of the block and $y_{1,t}\in\{0,1\}$ is the signal of player 1's action in the t^{th} period of the block.[28]) If $E=0$, the players "pass the review" and (C,L) is prescribed in the next block. If $E=1$, then with some probability $q\in[0,1]$ (which also remains to be determined), the players fail the review and (D,R) is prescribed forever. With the complementary probability $1-q$, the players pass the review anyway and (C,L) is prescribed in the next block.

We show that there exists $c>0$ and $\bar{\delta}<1$ such that, for any $\delta>\bar{\delta}$, the parameters ρ and q can be chosen so that this strategy profile is an equilibrium that yields payoff $v>2-c\left(\left(1-\delta\right)/\left(-\log\left(1-\delta\right)\right)\right)^{1/2}$ for each player.[29]

[28] The event E does not depend on the signals $\left(y_{2,t}\right)_{t=1}^{T}$ of player 2's action. Indeed, monitoring player 2 is unnecessary, as player 2 takes a static best response at action profile (C,L).

[29] We write v instead of v_i here, since the players' payoffs are the same.

Let p be the probability that $E = 1$ when player 1 takes C throughout a block; let p_1 be the probability that $E = 1$ when player 1 takes D once and takes C $T-1$ times; and let p_T be the probability that $E = 1$ when player 1 takes D throughout. Observe that v is given by

$$v = \left(1-\delta^T\right) 2 + \delta^T \left(1 - pq\right) v \quad \Longleftrightarrow \quad v = 2 - \frac{\delta^T pqv}{1-\delta^T}. \tag{PK}$$

At the same time, the incentive conditions that player 1 prefers to take C throughout a block where (C, L) is prescribed, rather than taking D in period 1 only, or always taking D, are

$$1 - \delta \leq \delta^T \left(p_1 - p\right) qv \qquad \text{and} \tag{IC$_1$}$$

$$1 - \delta^T \leq \delta^T \left(p_T - p\right) qv. \tag{IC$_T$}$$

Conditions (IC$_1$) and (IC$_T$) are obviously necessary for the review strategy to be an equilibrium; moreover, as shown by Matsushima (2004, p. 846), they are also sufficient.[30] It thus suffices to find $c > 0$, $\bar{\delta} < 1$, ρ, and q such that, for any $\delta > \bar{\delta}$, (PK), (IC$_1$), and (IC$_T$) hold, and $v > 2 - c\left(\left(1-\delta\right)/\left(-\log\left(1-\delta\right)\right)\right)^{1/2}$.

Define

$$v^* = 2 - \frac{1-\delta}{1-\delta^T}\frac{p}{p_1 - p} \qquad \text{and} \qquad q = \left(\frac{2\delta^T}{1-\delta}\left(p_1 - p\right) - \frac{\delta^T}{1-\delta^T} p\right)^{-1}. \tag{21}$$

With these definitions, (PK) and (IC$_1$) hold with equality, with $v = v^*$. We show that

$$\lim_{\delta \to 1} \frac{\frac{1-\delta}{1-\delta^T}\frac{p}{p_1-p}}{\sqrt{\frac{1-\delta}{-\log(1-\delta)}}} < \frac{5\sqrt{\rho}e^\rho}{e^\rho - 1} \text{ for all } \rho > 0, \tag{22}$$

$$\lim_{\rho \to 0}\lim_{\delta \to 1} \frac{2\delta^T\left(p_1 - p\right)}{1-\delta} - \frac{\delta^T p}{1-\delta^T} > 1, \qquad \text{and} \tag{23}$$

$$\lim_{\rho \to 0}\lim_{\delta \to 1} \frac{\delta^T q\left(p_T - p\right)}{1-\delta^T} > 1. \tag{24}$$

Given these inequalities, the proof is completed by first taking $\rho > 0$ and $\bar{\delta}_1 > 0$ such that the inequalities in (23) and (24) hold for ρ and all $\delta > \bar{\delta}_1$, then taking $\bar{\delta}_2 > 0$ such that

[30] Matsushima considered repeated games with two players and conditionally independent signals. Conditional independence implies that a player does not learn about her opponents signals during a review block, just as players do not learn about the mediator's signals in Γ^B. The same argument thus applies here.

the inequality in (22) holds for ρ for all $\delta > \bar{\delta}_2$, and finally taking $c = 5\sqrt{\rho}e^{\rho}/(e^{\rho}-1)$ and $\bar{\delta} = \max\{\bar{\delta}_1, \bar{\delta}_2\}$.

We now establish (22)–(24). Let $k = \left\lfloor \left(\sqrt{T}/2\right)\left(\sqrt{T} - \sqrt{-\log(1-\delta)}\right)\right\rfloor$. Note that

$$\begin{aligned} p &= \Pr\left(\sum_{t=2}^{T} y_{1,t} < k\right) + \frac{1}{2}\Pr\left(\sum_{t=2}^{T} y_{1,t} = k\right) < \Pr\left(\sum_{t=2}^{T} y_{1,t} \leq k\right), \qquad \text{and} \\ p_1 - p &= \frac{1}{4}\Pr\left(\sum_{t=2}^{T} y_{1,t} = k\right) = \frac{(T-1)!}{k!\,(T-1-k)!}\left(\frac{1}{2}\right)^{T+1} \geq \frac{(T)!}{k!\,(T-k)!}\left(\frac{1}{2}\right)^{T+2}, \end{aligned} \tag{25}$$

where the last inequality holds because $k \leq T/2$.

We first establish (22). Recall that the $y_{1,t}$ are independent Bernoulli random variables. As shown by Zhu, Li, and Hayashi (2022, Theorem 2.1),

$$\frac{\Pr\left(\sum_{t=2}^{T} y_{1,t} \leq k\right)}{\Pr\left(\sum_{t=2}^{T} y_{1,t} = k\right)} \leq k + 1 - \frac{T}{2} + \sqrt{\left(k - 1 - \frac{T}{2}\right)^2 + 2k}.$$

Since

$$\frac{p}{p_1 - p} < 4\frac{\Pr\left(\sum_{t=2}^{T} y_{1,t} \leq k\right)}{\Pr\left(\sum_{t=2}^{T} y_{1,t} = k\right)} \qquad \text{and} \qquad \lim_{\delta\to 1} \frac{\frac{1-\delta}{1-\delta^T}\left(k + 1 - \frac{T}{2} + \sqrt{\left(k - 1 - \frac{T}{2}\right)^2 + 2k}\right)}{\sqrt{\frac{1-\delta}{-\log(1-\delta)}}} = \frac{\sqrt{\rho}e^{\rho}}{e^{\rho}-1},$$

where the second line follows by l'Hopital's rule, we have

$$\lim_{\delta\to 1} \frac{\frac{1-\delta}{1-\delta^T}\frac{p}{p_1-p}}{\sqrt{\frac{1-\delta}{-\log(1-\delta)}}} \leq \lim_{\delta\to 1} \frac{\frac{1-\delta}{1-\delta^T}4\frac{\Pr\left(\sum_{t=2}^{T} y_{1,t}\leq k\right)}{\Pr\left(\sum_{t=2}^{T} y_{1,t}=k\right)}}{\sqrt{\frac{1-\delta}{-\log(1-\delta)}}} = \frac{4\sqrt{\rho}e^{\rho}}{e^{\rho}-1} < \frac{5\sqrt{\rho}e^{\rho}}{e^{\rho}-1},$$

which establishes (22).

We next establish (23). Applying Stirling's formula to (25), we have

$$p_1 - p \geq \frac{\sqrt{2\pi(T-1)}}{4e^2\sqrt{k(T-1-k)}}\left(\frac{T-1}{2k}\right)^k \left(\frac{T-1}{2(T-1-k)}\right)^{T-1-k}. \tag{26}$$

Therefore,

$$\lim_{\rho\to 0}\lim_{\delta\to 1}\frac{\delta^T(p_1-p)}{1-\delta}\geq\lim_{\rho\to 0}\lim_{\delta\to 1}\frac{\delta^T}{1-\delta}\frac{\sqrt{2\pi(T-1)}}{4e^2\sqrt{k(T-1-k)}}\left(\frac{T-1}{2k}\right)^k\left(\frac{T-1}{2(T-1-k)}\right)^{T-1-k}=\infty.$$

On the other hand, $\lim_{\rho\to 0}\lim_{\delta\to 1}\frac{\delta^T p}{1-\delta^T}\leq\lim_{\delta\to 1}\frac{\delta^T}{1-\delta^T}=\lim_{\rho\to 0}\frac{e^{-\rho}}{1-e^{-\rho}}=\infty$, which establishes (23).

Finally, we establish (24). We will show that $\lim_{\delta\to 1}p=0$ and $\lim_{\delta\to 1}p_T=1$. Hence, for sufficiently large δ, $p_T-p\geq 1/2$. This implies (24), as we have

$$\begin{aligned}
&\lim_{\rho\to 0}\lim_{\delta\to 1}\frac{\delta^T q(p_T-p)}{1-\delta^T}\\
=\ &\lim_{\rho\to 0}\lim_{\delta\to 1}\frac{\delta^T}{1-\delta^T}\frac{1}{\delta^T}\frac{p_T-p}{2\frac{p_1-p}{1-\delta}-\frac{p}{1-\delta^T}} \qquad \text{by (21)}\\
\geq\ &\lim_{\rho\to 0}\lim_{\delta\to 1}\frac{1}{1-\delta^T}\frac{\frac{1}{2}}{2\frac{p_1-p}{1-\delta}}\\
\geq\ &\lim_{\rho\to 0}\lim_{\delta\to 1}\frac{1-\delta}{1-\delta^T}\left(\frac{\sqrt{e(T-1)}}{2\pi\sqrt{k(T-1-k)}}\left(\frac{T-1}{2k}\right)^k\left(\frac{T-1}{2(T-1-k)}\right)^{T-1-k}\right)^{-1}=\infty,
\end{aligned}$$

where the second inequality follows by applying Stirling's formula to (25).

It remains to show that $\lim_{\delta\to 1}p=0$ and $\lim_{\delta\to 1}p_T=1$. Note that the random variable $2y_{1,t}-1$ has zero mean and unit variance when player 1 takes C. Thus, by the Berry–Esseen theorem, there exists an absolute constant C_0 such that

$$\begin{aligned}
p\ &=\ \Pr^{\text{player 1 takes }C}\left(\frac{1}{\sqrt{T}}\sum_{t=1}^{T}(2y_{1,t}-1)\leq-\sqrt{-\log(1-\delta)}\right)\\
&\leq\ \Phi\left(-\sqrt{-\log(1-\delta)}\right)+C_0\left(\frac{\mathbb{E}^{\text{player 1 takes }C}\left[|2y_{1,t}-1|^3\right]}{\sqrt{T}}\right)\xrightarrow{\delta\to 1}0.
\end{aligned}$$

On the other hand, $(4y_{1,t}-1)/\sqrt{3}$ has zero mean and unit variance when player 1 takes D.

Thus, again by Berry–Esseen,

$$\begin{aligned} p_T &= \Pr^{\text{player 1 takes } D}\left(\frac{1}{\sqrt{T}}\sum_{t=1}^{T}(2y_{1,t}-1) \leq -\sqrt{-\log(1-\delta)}\right) \\ &= \Pr^{\text{player 1 takes } D}\left(\frac{1}{\sqrt{T}}\sum_{t=1}^{T}\frac{4y_{1,t}-1}{\sqrt{3}} \leq \frac{\sqrt{T}-2\sqrt{-\log(1-\delta)}}{\sqrt{3}}\right) \\ &\geq \Phi\left(\frac{\sqrt{T}-2\sqrt{-\log(1-\delta)}}{\sqrt{3}}\right) + C_0\left(\frac{\mathbb{E}^{\text{player 1 takes } D}\left[\left|(4y_{1,t}-1)/\sqrt{3}\right|^3\right]}{\sqrt{T}}\right) \xrightarrow{\delta\to 1} 1, \end{aligned}$$

completing the proof.

C Proof of Theorem 3

We first show that first-order inefficiency in the blind game is no less than (14). Fix $\delta < 1$, $w \in (0, \bar{u})$, and a Nash equilibrium in the blind game where the agent's payoff is w. Let $\mu \in \Delta((A \times Y)^\infty)$ and $\alpha \in \Delta(A)$ be the corresponding outcome and occupation measure. Let $\hat{u}_t = u(c_t) - \psi(a_t) - w$.

By feasibility, the principal's payoff is at most

$$(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\mathbb{E}^{\mu}\left[\bar{F}(w+\hat{u}_t)\right].$$

At the same time, incentive compatibility implies inequality (15) (see Lemma 22 in the online appendix), and promise keeping implies

$$(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\mathbb{E}^{\mu}[\hat{u}_t] = 0. \tag{27}$$

The following is the key lemma.

Lemma 10 *There exist $c, \varepsilon > 0$ such that, for any sufficiently large $\delta < 1$, we have*

$$\max_{\substack{(\hat{u}_t(h))_{t,h} \geq -\left(w+h\left(\bar{A}\right)\right) \\ s.t.\ (15)\ and\ (27)}} (1-\delta) \sum_{t=1}^{\infty} \delta^{t-1} \mathbb{E}^{\mu} \left[\bar{F}\left(w + \hat{u}_t(h)\right) \right]$$
$$\leq \bar{F}(w) + \frac{1-\delta}{\delta} \mathbb{E}^{\alpha} \left[\frac{\psi'(a)^2}{\mathcal{I}(a)} \right] \frac{\bar{F}''(w)}{2} + c(1-\delta)^{1+\varepsilon}. \tag{28}$$

We sketch the proof of Lemma 10, relegating the details to the online appendix. Subtracting $\bar{F}(w)$ from both sides of (28), dividing both sides by $\frac{1-\delta}{\delta} \frac{\bar{F}''(w)}{2}$, and taking a second-order Taylor approximation of the LHS (where the first-order term is zero by (27)) gives

$$\inf_{(\hat{u}_t(h))_{t,h} \geq -\left(w+\psi\left(\bar{A}\right)\right) \text{ s.t. } (15)} \sum_{t=1}^{\infty} \delta^t \mathbb{E}^{\mu} \left[\hat{u}_t(h)^2 \right].$$

To establish Lemma 10, it suffices to show that the value of this program exceeds $\mathbb{E}^{\alpha}\left[\psi'(a)^2 / \mathcal{I}(a)\right]$. To see this, take a common Lagrange multiplier of $2(1-\delta)\delta^{t-1}$ on (15) for each t. Then, by weak duality, the value of the program is no less than

$$\inf_{(\hat{u}_t(h))_{t,h}} \sum_{t=1}^{\infty} \mathbb{E}^{\mu} \left[\delta^t \hat{u}_t(h)^2 - 2\frac{\psi'(a_t)\,\nu(y_t|a_t)}{\mathcal{I}(a_t)} \frac{1-\delta}{\delta} \sum_{t'=t+1}^{\infty} \delta^{t'} \hat{u}_{t'}(h) \right] + 2\mathbb{E}^{\alpha} \left[\frac{\psi'(a)^2}{\mathcal{I}(a)} \right]. \tag{29}$$

Taking the first-order condition for $\hat{u}_t(h)$ and substituting into (29) as in the text gives

$$-\left(\frac{1-\delta}{\delta}\right)^2 \sum_{t=2}^{\infty} \delta^t \mathbb{E}^{\mu} \left[\left(\sum_{t'=1}^{t-1} \frac{\psi'(a_{t'})\,\nu(y_{t'}|a_{t'})}{\mathcal{I}(a_{t'})} \right)^2 \right] + 2\mathbb{E}^{\alpha} \left[\frac{\psi'(a)^2}{\mathcal{I}(a)} \right] = \mathbb{E}^{\alpha} \left[\frac{\psi'(a)^2}{\mathcal{I}(a)} \right],$$

where the equality follows because, since $\left(\psi_t(a_t)\,v_t(y_t|a_t) / \mathcal{I}_t(a_t)\right)_t$ is a sequence of martingale increments, we have

$$\begin{aligned}
\mathbb{E}^{\mu} \left[\left(\sum_{t'=1}^{t-1} \frac{\psi'(a_{t'})\,\nu(y_{t'}|a_{t'})}{\mathcal{I}(a_{t'})} \right)^2 \right] &= \sum_{t'=1}^{t-1} \mathbb{E}^{\mu} \left[\left(\frac{\psi'(a_{t'})\,\nu(y_{t'}|a_{t'})}{\mathcal{I}(a_{t'})} \right)^2 \right] \\
&= \sum_{t'=1}^{t-1} \mathbb{E}^{\alpha_{t'}} \left[\left(\frac{\psi'(a)}{\mathcal{I}(a)} \right)^2 \mathbb{E}\left[\nu(y|a)^2\right] \right] = \sum_{t'=1}^{t-1} \mathbb{E}^{\alpha_{t'}} \left[\frac{\psi'(a)^2}{\mathcal{I}(a)} \right],
\end{aligned}$$

and hence

$$\left(\frac{1-\delta}{\delta}\right)^2 \sum_{t=2}^{\infty} \delta^t \mathbb{E}^{\mu}\left[\left(\sum_{t'=1}^{t-1} \frac{\psi'(a_{t'})\,\nu(y_{t'}|a_{t'})}{\mathcal{I}(a_{t'})}\right)^2\right] = \left(\frac{1-\delta}{\delta}\right)^2 \sum_{t=2}^{\infty} \delta^t \sum_{t'=1}^{t-1} \mathbb{E}^{\alpha_{t'}}\left[\frac{\psi'(a)^2}{\mathcal{I}(a)}\right]$$

$$= (1-\delta)\sum_{t=1}^{\infty} \delta^{t-1}\mathbb{E}^{\alpha_t}\left[\frac{\psi'(a)^2}{\mathcal{I}(a)}\right] = \mathbb{E}^{\alpha}\left[\frac{\psi'(a)^2}{\mathcal{I}(a)}\right]. \tag{30}$$

At the same time, since $\bar{F}$ is strictly concave, there exists $\varepsilon_1 > 0$ such that the principal's payoff is at most $\bar{F}(w) - \varepsilon_1 \mathbb{E}^{\alpha}\left[(a - \bar{a}(w))^2\right]$. (See Lemma 19 in the online appendix.) So, together with Lemma 10, the following lemma establishes that first-order inefficiency in the blind game is no less than (14).

Lemma 11 *There exist $\hat{c}, \hat{\varepsilon} > 0$ such that, for any sufficiently large $\delta < 1$, we have*

$$\max_{\alpha} \min \left\{ \frac{1-\delta}{\delta}\mathbb{E}^{\alpha}\left[\frac{\psi'(a)^2}{\mathcal{I}(a)}\right] \frac{\bar{F}''(w)}{2} + c(1-\delta)^{1+\varepsilon}, -\varepsilon_1 \mathbb{E}^{\alpha}\left[(a-\bar{a}(w))^2\right] \right\}$$

$$\leq \frac{1-\delta}{\delta} \frac{\psi'(\bar{a}(w))^2}{\mathcal{I}(\bar{a}(w))} \frac{\bar{F}''(w)}{2} + \hat{c}(1-\delta)^{1+\hat{\varepsilon}}.$$

We now show that first-order inefficiency in the public game is no more than (14). The proof is constructive. As a first step, it is helpful to first construct a static contract that induces a target effort level $\bar{a} \in A$. In particular, if the agent is rewarded with a utility of $(\psi'(\bar{a})/\mathcal{I}(\bar{a}))\,\nu(y|\bar{a})$ following any signal realization y, she chooses a to maximize

$$\int_y \frac{\psi'(\bar{a})}{\mathcal{I}(\bar{a})}\nu(y|\bar{a})\,p(y|a)\,dy - \psi(a).$$

The solution is $a = \bar{a}$, because $a = \bar{a}$ satisfies the first-order condition

$$\int_y \frac{\psi'(\bar{a})}{\mathcal{I}(\bar{a})}\nu(y|\bar{a})\,p_a(y|a)\,dy = \psi'(a),$$

and the second-order condition holds by (11). Moreover, the expected reward equals zero.

Heuristically, the repeated game equilibrium is constructed by using the above reward to adjust the agent's continuation payoff w_t after each period t (so the agent's continuation payoff is a martingale), while targeting effort level $\bar{a}(w_t)$ in each period t. This heuristic requires two adjustments, however. First, if the score $\nu(y|a)$ is unbounded, we must truncate the reward for extreme scores, and then further adjust the reward so the agent's first-order condition is exactly satisfied. Second, it is convenient to target zero effort once the agent's continuation payoff w_t strays too far from its initial value w.

Formally, fix any $\varepsilon > 0$. Recursively, given the agent's promised continuation payoff $w_t(h^t)$ at history $h^t = (r_{t'}, y_{t'})_{t'=1}^{t-1}$, we define the recommended period t action $r_t(h^t)$ and consumption level $c_t(h^t)$ (which is independent of the period t signal y_t), as well as the next period continuation payoff $w_{t+1}(h^t, y_t)$, as follows. First, say that a history h^t is *regular* if $|w_t(h^t) - w| \leq (1-\delta)^{1/2-\varepsilon}$, and *irregular* otherwise. At a regular history, define $r_t(h^t) = \bar{a}(w_t(h^t))$, $c_t(h^t) = \bar{c}(w_t(h^t))$, and

$$w_{t+1}\left(h^t, y_t\right) = w_t\left(h^t\right) + \frac{1-\delta}{\delta} x_{r_t(h^t)}(y),$$

where, for each $\bar{a} \in A$, $x_{\bar{a}}(y)$ is an adjusted version of $(\psi'(\bar{a})/\mathcal{I}(\bar{a}))\,\nu(y|\bar{a})$ (constructed in Lemma 25 in the online appendix) that satisfies

$$\begin{aligned}
\bar{a} &\in \underset{a}{\operatorname{argmax}} \int_y x_{\bar{a}}(y)\, p(y|a)\, dy - \psi(a), \\
\int_y x_{\bar{a}}(y)\, p(y|\bar{a})\, dy &= 0, && (31)\\
\int_y x_{\bar{a}}(y)^2\, p(y|\bar{a})\, dy &= \frac{\psi'(\bar{a})^2}{\mathcal{I}(\bar{a})} + O(1-\delta), \qquad \text{and} && (32)\\
|x_{\bar{a}}(y)| &\leq (1-\delta)^{-1/4}.
\end{aligned}$$

At an irregular history, define $r_t(h^t) = 0$, $c_t(h^t) = u^{-1}(w_t(h^t))$, and $w_{t+1}(h^t, y_t) = w_t(h^t)$ for all y_t. Note that the initial history h^1 is regular, and that if a history h^t is irregular, then so is every subsequent history. Note also that, by construction, $|w_{t+1}(h^t, y_t) - w_t(h^t)| = O(1-\delta)^{3/4}$ for every regular history h^t and signal y_t. Since $|w_t(h^t) - w| \leq (1-\delta)^{1/2-\varepsilon}$ for every regular history h^t, this implies that, for sufficiently large $\delta < 1$, we have $w_t(h^t) \in [0, 2w]$ for *every* history h^t.

The proof is completed by the following lemma, which shows that the first-order inefficiency of this equilibrium is no more than (14).

Lemma 12 *There exist $\tilde{c}, \tilde{\varepsilon} > 0$ such that, for any sufficiently large $\delta < 1$, the principal's payoff in the above equilibrium is no less than*

$$\bar{F}(w) + \frac{1-\delta}{\delta} \frac{\psi'(\bar{a}(w))^2}{\mathcal{I}(\bar{a}(w))} \frac{\bar{F}''(w)}{2} - \tilde{c}(1-\delta)^{1+\tilde{\varepsilon}}.$$

Intuitively, since $w_t(h^t)$ is a martingale with volatility of order $(1-\delta)^2$ (by (31) and (32)), it is very unlikely that $w_t(h^t)$ moves more than a $O\left((1-\delta)^{1/2}\right)$ distance away from w within a timeframe that has more than an $O(1-\delta)$ payoff impact. Consequently, the principal's payoff is almost entirely determined by her payoff at regular histories, and thus equals

$$\mathbb{E}^{\mu}\left[(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\bar{F}\left(w_t\left(h^t\right)\right)\right] + o(1-\delta),$$

where μ is the equilibrium outcome. Taking a second-order Taylor expansion around $w_t(h^t) = w$ and ignoring the remainder, this equals

$$\bar{F}(w) + \mathbb{E}^{\mu}\left[(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\frac{\bar{F}''(w)}{2}\left(w_t\left(h^t\right)-w\right)^2\right].$$

Since

$$w_t\left(h^t\right) = w + \frac{1-\delta}{\delta}\sum_{t'=1}^{t-1} x_{a_{t'}}(y_{t'}) \qquad \text{for all regular histories } h^t$$

and (32) holds, the same calculation as for (30) gives

$$\mathbb{E}^{\mu}\left[(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\frac{\bar{F}''(w)}{2}\left(w_t\left(h^t\right)-w\right)^2\right] = \frac{1-\delta}{\delta}\mathbb{E}^{\alpha}\left[\frac{\psi'(a)}{\mathcal{I}(a)}\right]\frac{\bar{F}''(w)}{2} + o(1-\delta),$$

where α is the equilibrium occupation measure. Finally, since $w_t(h^t)$ is close to w with high probability under α, we have $\mathbb{E}^{\alpha}[\psi'(a)/\mathcal{I}(a)] = \psi'(\bar{a}(w))/\mathcal{I}(\bar{a}(w)) + o(1)$, completing the proof.

References

[1] Abreu, Dilip, David Pearce, and Ennio Stacchetti. "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring." *Econometrica* 58.5 (1990): 1041-1063.

[2] Abreu, Dilip, Paul Milgrom, and David Pearce. "Information and Timing in Repeated Partnerships." *Econometrica* 59.6 (1991): 1713-1733.

[3] Aoyagi, Masaki. "Information Feedback in a Dynamic Tournament." *Games and Economic Behavior* 70.2 (2010): 242-260.

[4] Athey, Susan, and Kyle Bagwell. "Optimal Collusion with Private Information." *RAND Journal of Economics* 32.3 (2001): 428-465.

[5] Ball, Ian. "Dynamic Information Provision: Rewarding the Past and Guiding the Future." *Econometrica*, 91.4 (2023): 1363-1391.

[6] Buldygin, Valerii and Yu Kozachenko, *Metric Characterization of Random Variables and Random Processes*, Vol. 188, American Mathematical Society (2000).

[7] Ederer, Florian. "Feedback and Motivation in Dynamic Tournaments." *Journal of Economics & Management Strategy* 19.3 (2010): 733-769.

[8] Ely, Jeffrey, George Georgiadis, and Luis Rayo, "Feedback Design in Dynamic Moral Hazard," *Working Paper* (2024).

[9] Ely, Jeffrey, George Georgiadis, Sina Khorasani, and Luis Rayo, "Optimal Feedback in Contests." *Review of Economic Studies*, 90.5 (2023): 2370-2394.

[10] Ely, Jeffrey C., and Martin Szydlowski. "Moving the Goalposts." *Journal of Political Economy* 128.2 (2020): 468-506.

[11] Forges, Francoise. "An Approach to Communication Equilibria." *Econometrica* 54.6 (1986): 1375-1385.

[12] Frick, Mira, Ryota Iijma, and Yuhta Ishii. "Monitoring with Rich Data." *Working Paper* (2024).

[13] Fuchs, William. "Contracting with Repeated Moral Hazard and Private Evaluations." *American Economic Review* 97.4 (2007): 1432-1448.

[14] Fudenberg, Drew, and David Levine. "Continuous Time Limits of Repeated Games with Imperfect Public Monitoring." *Review of Economic Dynamics* 10.2 (2007): 173-192.

[15] Fudenberg, Drew, David Levine, and Eric Maskin. "The Folk Theorem with Imperfect Public Information." *Econometrica* 62.5 (1994): 997-1039.

[16] Gershkov, Alex, and Motty Perry. "Tournaments with Midterm Reviews." *Games and Economic Behavior* 66.1 (2009): 162-190.

[17] Goldlücke, Susanne, and Sebastian Kranz. "Infinitely Repeated Games with Public Monitoring and Monetary Transfers." *Journal of Economic Theory* 147.3 (2012): 1191-1221.

[18] Green, Edward J., and Robert H. Porter. "Noncooperative Collusion under Imperfect Price Information." *Econometrica* 52.1 (1984): 87-100.

[19] Halac, Marina, Navin Kartik, and Qingmin Liu. "Contests for Experimentation." *Journal of Political Economy* 125.5 (2017): 1523-1569.

[20] Holmström, Bengt, and Paul Milgrom. "Aggregation and Linearity in the Provision of Intertemporal Incentives," *Econometrica* 55.2 (1987): 303-328.

[21] Hörner, Johannes, and Satoru Takahashi. "How Fast do Equilibrium Payoff Sets Converge in Repeated Games?" *Journal of Economic Theory* 165 (2016): 332-359.

[22] Kandori, Michihiro. "Introduction to Repeated Games with Private Monitoring." *Journal of Economic Theory* 102.1 (2002): 1-15.

[23] Kandori, Michihiro, and Hitoshi Matsushima. "Private Observation, Communication and Collusion." *Econometrica* 66.3 (1998): 627-652.

[24] Kandori, Michihiro, and Ichiro Obara. "Efficiency in Repeated Games Revisited: The role of Private Strategies." *Econometrica* 74.2 (2006): 499-519.

[25] Levin, Jonathan. "Relational Incentive Contracts." *American Economic Review* 93.3 (2003): 835-857.

[26] Lizzeri, Alessandro, Margaret Meyer, and Nicola Persico. "The Incentive Effects of Interim Performance Evaluations." *Working Paper* (2002).

[27] Madrigal, Vicente. "On the Non-existence of Efficient Equilibria of Repeated Principal Agent Games with Discounting." *Working Paper* (1986).

[28] Matsushima, Hitoshi. "Multimarket Contact, Imperfect Monitoring, and Implicit Collusion." *Journal of Economic Theory* 98.1 (2001): 158-178.

[29] Matsushima, Hitoshi. "Repeated Games with Private Monitoring: Two Players." *Econometrica* 72.3 (2004): 823-852.

[30] Mirrlees, James A. "The Theory of Moral Hazard and Unobservable Behaviour: Part I." *Working Paper* (1975) (published in *Review of Economic Studies* 66.1 (1999): 3-21).

[31] Orlov, Dmitry, Andrzej Skrzypacz, and Pavel Zryumov. "Persuading the Principal to Wait." *Journal of Political Economy* 128.7 (2020): 2542-2578.

[32] Radner, Roy. "Repeated Principal-Agent Games with Discounting," *Econometrica* 53.5 (1985): 1173-1198.

[33] Rahman, David. "The Power of Communication." *American Economic Review* 104.11 (2014): 3737-3751.

[34] Rubinstein, Ariel, "An Optimal Conviction Policy for Offenses that May have been Committed by Accident," in *Applied Game Theory*, ed. by Brams, Schotter, and Schwodiauer, (1979) 406–413, Physical-Verlag: Heidleberg.

[35] Rubinstein, Ariel, and Menahem E. Yaari. "Repeated Insurance Contracts and Moral Jazard." *Journal of Economic Theory* 30.1 (1983): 74-97.

[36] Sadzik, Tomasz, and Ennio Stacchetti. "Agency Models with Frequent Actions." *Econometrica* 83.1 (2015): 193-237.

[37] Sannikov, Yuliy. "Games with Imperfectly Observable Actions in Continuous Time." *Econometrica* 75.5 (2007): 1285-1329.

[38] Sannikov, Yuliy. "A Continuous-Time Version of the Principal-Agent Problem." *Review of Economic Studies* 75.3 (2008): 957-984.

[39] Sannikov, Yuliy, and Andrzej Skrzypacz. "Impossibility of Collusion under Imperfect Monitoring with Flexible Production." *American Economic Review* 97.5 (2007): 1794-1823.

[40] Sannikov, Yuliy, and Andrzej Skrzypacz. "The Role of Information in Repeated Games with Frequent Actions." *Econometrica* 78.3 (2010): 847-882.

[41] Smolin, Alex. "Dynamic Evaluation Design." *American Economic Journal: Microeconomics* 13.4 (2021): 300-331.

[42] Spear, Stephen E., and Sanjay Srivastava. "On Repeated Moral Hazard with Discounting." *Review of Economic Studies* 54.4 (1987): 599-617.

[43] Sugaya, Takuo. "Folk Theorem in Repeated Games with Private Monitoring." *Review of Economic Studies* 89.4 (2022): 2201-2256.

[44] Sugaya, Takuo, and Alexander Wolitzky. "Bounding Equilibrium Payoffs in Repeated Games with Private Monitoring." *Theoretical Economics* 12 (2017): 691-729.

[45] Sugaya, Takuo, and Alexander Wolitzky. "Maintaining Privacy in Cartels." *Journal of Political Economy* 126.6 (2018): 2569-2607.

[46] Sugaya, Takuo, and Alexander Wolitzky. "Monitoring vs. Discounting in Repeated Games." *Econometrica*, Forthcoming (2023).

[47] Wood, David C. "The Computation of Polylogarithms." *Working Paper* (1992).

[48] Zhu, Huangjun, Zihao Li, and Masahito Hayashi. "Nearly Tight Universal Bounds for the Binomial Tail Probabilities." *Working Paper* (2022).

Online Appendix

D Proof of Lemma 2

The proof is similar to (but simpler than) the proof of Lemma 6 of Sugaya and Wolitzky (2023). To show that B is self-generating, it suffices to show that the extreme points of any ball $B' \subseteq B$ of radius $r/2$ are decomposable on B'.

Lemma 13 (Sugaya and Wolitzky (2023), Lemma 10) *Suppose that for any ball $B' \subseteq B$ with radius $r/2$ and any direction $\lambda \in \Lambda$, the point $\hat{v} = \operatorname{argmax}_{v' \in B'} \lambda \cdot v'$ is decomposable on B'. Then B is self-generating.*

We thus fix a ball $B' \subseteq B$ of radius $r/2$ and a direction $\lambda \in \Lambda$, and let $\hat{v} = \operatorname{argmax}_{v' \in B'} \lambda \cdot v'$. We construct (α, w) that decompose $\hat{v}$ on B'.

Since $k(\lambda, \bar{x}) \geq \max_{v' \in B} \lambda \cdot v' + \varepsilon$ by hypothesis, there exist α and $x : Y \to \mathbb{R}^N$ satisfying (IC), (HS$\bar{x}$), and

$$\lambda \cdot \left(u(\alpha) + \int_y x(y) \, p(y|\alpha) \, dy \right) \geq \max_{v' \in B} \lambda \cdot v' + \varepsilon/2 \geq \max_{v' \in B'} \lambda \cdot v' + \varepsilon/2. \tag{33}$$

To construct w, for each y, let

$$w(y) = \hat{v} + \frac{1-\delta}{\delta} \left(\hat{v} - u(\alpha) + x(y) - \int_{y'} p(y'|\alpha) \, x(y') \, dy' \right).$$

We show that (α, w) decomposes $\hat{v}$ on B' by verifying (PK), (IC), and (SG).

(PK): This holds by construction: we have $\int_y w(y) \, p(y|\alpha) \, dy = (1/\delta)(\hat{v} - (1-\delta) u(\alpha))$, and hence $(1-\delta) u(\alpha) + \delta \int_y w(y) \, p(y|\alpha) \, dy = \hat{v}$.

(IC): Setting aside the constant terms in $w(y)$, we see that an action a_i maximizes $(1-\delta) u_i(a_i, \alpha_{-i}) + \delta \int_y w_i(y) \, p(y|a_i, \alpha_{-i}) \, dy$ iff it maximizes $u_i(a_i, \alpha_{-i}) + \int_y x_i(y) \, p(y|a_i, \alpha_{-i}) \, dy$, which follows from (IC).

(SG): We start with a simple geometric observation.

Lemma 14 (Sugaya and Wolitzky (2023), Lemma 11) *For each $w \in \mathbb{R}^N$, we have $w \in B'$ if $\lambda \cdot (\hat{v} - w) \geq 0$ and*

$$d(\hat{v}, w) \leq \sqrt{(r/2) \, \lambda \cdot (\hat{v} - w)}. \tag{34}$$

We thus show that, for each y, $w(y)$ satisfies $\lambda \cdot (\hat{v} - w(y)) \geq 0$ and (34). Note that

$$\hat{v} - w(y) = \frac{1-\delta}{\delta}\left(u(\alpha) + \int_{y'} x(y')\, p(y'|\alpha)\, dy' - \hat{v} - x(y)\right).$$

By (HS$\bar{x}$) and (33), we have $\lambda \cdot (\hat{v} - w(y)) \geq (\delta/(1-\delta))\,\varepsilon/2$, and therefore

$$\sqrt{(r/2)\,\lambda \cdot (\hat{v} - w(y))} \geq \frac{1-\delta}{\delta}\sqrt{\frac{\delta}{1-\delta}\frac{\varepsilon r}{4}}. \tag{35}$$

Similarly, we have

$$\begin{aligned} d(\hat{v}, w) &\leq \frac{1-\delta}{\delta}\left(d(\hat{v}, u(\alpha)) + d\left(\int_{y'} x(y')\, p(y'|\alpha)\, dy', x(y)\right)\right) \\ &\leq \frac{1-\delta}{\delta}\left(\max_{u,u' \in F} \|u - u'\| + 2\bar{x}\right) \leq \frac{1-\delta}{\delta} 3\bar{x}. \end{aligned} \tag{36}$$

Comparing (35) and (36), we see that $w(y)$ satisfies (34) whenever $3\bar{x} \leq \sqrt{(\delta/(1-\delta))\,\varepsilon r/4}$, which holds by (9).

E Proof of Lemma 3

Since $\dim F^* = N$ and $v_i < \max_a u_i(a)$ for all i, there exist $\bar{\eta} > 0$ and $\bar{F} \subseteq F^*$ such that $\{w \in F^* : d(v, w) \leq \bar{\eta}\} \subseteq \bar{F}$, $\dim \bar{F} = N$, and $\underline{v}_i < w_i < \max_a u_i(a)$ for all i and $w \in \bar{F}$. Fix any such $(\bar{\eta}, \bar{F})$.

The following lemma is similar to Lemma 5 of Hörner and Takahashi (2016) or Lemma 7 and pp. 1750–1751 of Sugaya and Wolitzky (2023).

Lemma 15 *There exists $\bar{x} > \bar{u}$ such that $k(\lambda, \bar{x}) \geq \max_{v' \in \bar{F}} \lambda \cdot v'$ for all $\lambda \in \Lambda$.*

Proof. Let $\hat{x} > 0$ satisfy the conditions of Assumption 2. For each i, since $\underline{v}_i < w_i < \max_a u_i(a)$ for all $w \in \bar{F}$, there exist $\underline{\lambda}_i > -1$ and $\bar{\lambda}_i < 1$ such that (i) for all $\lambda \in \Lambda$ with $\lambda_i \leq \underline{\lambda}_i$, we have $\lambda \cdot u(\alpha) - \sum_{n \neq i} |\lambda_n|\, \hat{x} \geq \max_{w \in \bar{F}} \lambda \cdot w$ for all α satisfying $u_i(\alpha) = \underline{v}_i$; and (ii) for all $\lambda \in \Lambda$ with $\lambda_i \geq \bar{\lambda}_i$, we have $\lambda \cdot u(\alpha) - \sum_{n \neq i} |\lambda_n|\, \hat{x} \geq \max_{w \in \bar{F}} \lambda \cdot w$ for all α satisfying $u_i(\alpha) = \max_a u_i(a)$. Given such $(\underline{\lambda}_i, \bar{\lambda}_i)_i$, we define

$$\bar{x} = \sqrt{N}\left(2N + \max_i \frac{\max\{|\underline{\lambda}_i|, \bar{\lambda}_i\}}{\sqrt{(1 - \max\{|\underline{\lambda}_i|, \bar{\lambda}_i\}^2)/(N-1)}}\right)\hat{x}.$$

For each λ, we now construct $(\alpha, x(y))$ such that $\lambda \cdot (u(\alpha) + \mathbb{E}[x(y)|\alpha]) \geq \max_{w \in \bar{F}} \lambda \cdot w$ and (IC) and (HS$\bar{\text{x}}$) hold. To do so, fix any $i \in \text{argmax}\,|\lambda_i|$, and consider three cases.

(i) $\lambda_i \leq \underline{\lambda}_i$. In this case, take a minmax profile α^i and $(\hat{x}_j(y))_{j\neq i}$ that satisfy Assumption 2(i). Define $x_i(y) = \sum_{n\neq i} |\lambda_n| \hat{x}/\lambda_i$ and $x_j(y) = \hat{x}_j(y)$ for all y. (IC) holds for $j \neq i$ given Assumption 2(i), and (IC) holds for i since player i takes a best response in α^i and $x_i(y)$ is independent of y. (HS$\bar{\text{x}}$) holds with $\bar{x} \geq 2N\hat{x}$ given Assumption 2(i). Finally, $\lambda \cdot u(\alpha^i) + \mathbb{E}[\lambda \cdot x(y)] \geq \lambda \cdot u(\alpha^i) - \sum_{n\neq i} |\lambda_n| \hat{x} \geq \max_{w\in\bar{F}} \lambda \cdot w$.

(ii) $\lambda_i \geq \underline{\lambda}_i$. In this case, take $\bar{a}^i \in \max u_i(a)$ and $(\hat{x}_j(y))_{j\neq i}$ that satisfies (6) for player j. Define $x_i(y) = \sum_{n\neq i} |\lambda_n| \hat{x}/\lambda_i$ and $x_j(y) = \hat{x}_j(y)$ for all y. The argument is now the same as case (ii).

(iii) Otherwise, there exists $n \neq i$ such that $|\lambda_i| / |\lambda_n| \leq \max\{|\underline{\lambda}_i|, \bar{\lambda}_i\}/\sqrt{(1-\max\{|\underline{\lambda}_i|, \bar{\lambda}_i\}^2)/(N-1)}$. Fix such n. Next, fix $a \in A$ such that $\lambda \cdot u(a) \geq \max_{w\in\bar{F}} \lambda \cdot w$. For player i, take $x^i(y)$ such that Assumption 2(ii) holds for a, (i,n), and $c = \text{sign}(-\lambda_n/\lambda_i)$. For player $j \neq i$, take $x^j(y)$ such that Assumption 2(ii) holds for a, (j,i), and $c = \text{sign}(-\lambda_j/\lambda_i)$. We then define $x_i(y) = x^i(y) - \sum_{j\neq i} \frac{\lambda_j}{\lambda_i} x^j(y)$, $x_n(y) = x^n(y) - \frac{\lambda_i}{\lambda_n} x^i(y)$, and $x_j(y) = x^j(y)$ for $j \neq i, n$. Then, (IC) holds for player i since $a_i \in \arg\max_{a_i'} u(a_i', a_{-i}) + \mathbb{E}[x^i(y)|a_i', a_{-i}]$ and $a_i \in \arg\max_{a_i'} \mathbb{E}[-(\lambda_j/\lambda_i) x^j(y)|a_i', a_{-i}]$ for all $j \neq i$ by Assumption 2(ii). (IC) holds for player n since $a_n \in \arg\max_{a_n'} u(a_n', a_{-n}) + \mathbb{E}[x^n(y)|a_n', a_{-n}]$ and $a_n \in \arg\max_{a_n'} \mathbb{E}[-(\lambda_i/\lambda_n) x^i(y)|a_n', a_{-n}]$ by Assumption 2(ii). In addition, (IC) holds for player $j \neq i, n$ since $a_j \in \arg\max_{a_j'} u(a_j', a_{-j}) + \mathbb{E}\left[x^j(y)|a_j', a_{-j}\right]$ by Assumption 2(ii). Finally, (HS$\bar{\text{x}}$) holds since $\lambda \cdot x(y) = 0$ for all y and $\|x(y)\| \leq \sqrt{N} \sum_j |x_j(y)| \leq \sqrt{N}(2N + |\lambda_i/\lambda_n|)\hat{x} \leq \bar{x}$. ■

By Lemma 15, it suffices to find $c > 0$ and $\bar{\delta} < 1$ such that, for all $\delta > \bar{\delta}$, there exist $\varepsilon > 0$ and a ball B with radius $r > 0$ such that

$$\max_{v'\in\bar{F}} \lambda' \cdot v' \geq \max_{v'\in B} \lambda' \cdot v' + \varepsilon \quad \text{for all } \lambda' \in \Lambda, \tag{37}$$

$$r\varepsilon \geq 36\bar{x}^2(1-\delta), \quad \text{and} \tag{38}$$

$$d(v, B) \leq c(1-\delta)^{\beta^*/2}. \tag{39}$$

If $\beta^* = 1$ then, as in Lemma 3 of Hörner and Takahashi (2016), it suffices to take any $o \in \text{int}\left(\bar{F}\right)$ and any $\ell > 0$ sufficiently large compared to $36\bar{x}^2$, let $r = (1-\delta)^{1/2}$, and take B to have radius r and center $(1-\ell r)v + \ell r o$.

For the rest of the proof, we assume that $\beta^* > 1$. We first derive a geometric condition for $w \in F^*$, similar to Lemma 14.

Lemma 16 *There exist $\lambda \in \Lambda_v^*$, $\rho > 0$, and $\kappa > 0$ such that, if $d(v, w) < \rho$ and $\kappa d(v, w)^{\beta^*} < \lambda \cdot (v - w)$, then $w \in \text{int}(F^*)$.*

Proof. Since F^* is full-dimensional and has min-curvature of order at most β at v, there exist $\bar{\varepsilon} > 0$ and $\kappa > 0$ such that, for all $w \in \text{bnd}(F^*)$ satisfying $d(v, w) < \bar{\varepsilon}$, we have $\lambda \cdot (v - w) <$

$\kappa d(v,w)^{\beta} \leq \kappa d(v,w)^{\beta^*}$ for all $\lambda \in \Lambda_v^*$. Let $B_{\varepsilon'}(v) = \{w \in \mathbb{R}^N : d(v,w) = \varepsilon'\}$. Since F^* is full-dimensional, there exists $\lambda \in \Lambda_v^*$, $\varepsilon' > 0$, and $t > 0$ such that $C := B_{\varepsilon'}(v) - t\lambda \subseteq F^*$. Fix such λ, ε', and t, and let $\varepsilon = \min\{\bar{\varepsilon}, \varepsilon', t\}$.

Now fix any $\rho < \min\{\varepsilon, (t/2\kappa)^{1/\beta^*}\}$ and $d < \rho$, and let $G = \{w \in B_d(v) : \lambda \cdot (v-w) \geq 2\kappa d^{\beta^*}\}$. We wish to show that $G \subseteq F^*$ (and in particular $G \subseteq \text{int}(F^*)$, since $G \cap \text{bnd}(F^*) = \emptyset$).

To see this, let $W = B_d(v) \cap \text{bnd}(F^*)$, $H = \{w : \lambda \cdot (v-w) = \kappa d^{\beta^*}\}$, $H' = \{w : \lambda \cdot (v-w) = t\}$, and $D = C \cap H'$. Since $d < \rho < \min\{\varepsilon, (t/\kappa)^{1/\beta^*}\}$, G lies in between H and H'. In addition, the projection of G onto H is a subset of the projection of W onto H, and the projection of G onto H' is a subset of D. Hence, we have $G \subseteq \text{co}(W \cup D)$. Finally, since $W \subseteq F^*$ and $D \subseteq C \subseteq F^*$, and F^* is convex, we have $\text{co}(W \cup D) \subseteq F^*$, so $G \subseteq F^*$. ■

Lemma 17 *There exist $\bar{c} > 0$, $\eta > 0$, and $\bar{\delta} < 1$ such that, for all $\delta > \bar{\delta}$, there exists a ball $B \subseteq \bar{F}$ of radius $r = \eta(1-\delta)^{1-\beta^*/2}$ satisfying $d(v,B) = \bar{c}(1-\delta)^{\beta^*/2}$.*

Proof. Fix $\lambda \in \Lambda_v^*$, $\rho > 0$, and $\kappa > 0$ as in Lemma 16. Given $\bar{c}$ and η to be determined, let B be the ball with radius $r = \eta(1-\delta)^{1-\beta^*/2}$ and center $o = v - (r+d)\lambda$, where $d = \bar{c}(1-\delta)^{\beta^*/2}$, and take any $\hat{w} \in \partial B$. Let $x = \lambda \cdot (\hat{w} - o)$, so that $x\lambda$ is the projection of $\hat{w} - o$ on λ. Then,

$$\begin{aligned} \|v - \hat{w}\|^2 &= \|v - o - x\lambda\|^2 + \|\hat{w} - o - x\lambda\|^2 = (r+d-x)^2 + r^2 - x^2, \qquad \text{and} \\ \lambda \cdot (v - \hat{w}) &= r + d - x. \end{aligned}$$

Recall that, by construction, $\{w \in F^* : d(v,w) \leq \bar{\eta}\} \subseteq \bar{F}$. Since $d(v,w) \leq d(v,o) + d(o,w) \leq 2r + d$ for all $w \in B$, it suffices to show that $2r + d \leq \bar{\eta}$ and $B \subseteq F^*$. By Lemma 16, if $d(v,w) < \rho$ and $\kappa d(v,w)^{\beta^*} \leq \lambda \cdot (v-w)$ then $w \in F^*$. Since $x \in [-r,r]$, it suffices to find $\bar{c}$, η, and $\bar{\delta}$ such that, for all $\delta > \bar{\delta}$, we have

$$\begin{aligned} 2r + d &\leq \bar{\eta}, & & (40) \\ ((r+d) - x)^2 + r^2 - x^2 &\leq \rho^2 \qquad \text{for all } x \in [-r,r], \qquad \text{and} & & (41) \\ \max_{x \in [-r,r]} f(x,\delta,\beta^*) &\leq 0, & & (42) \end{aligned}$$

where

$$f(x,\delta,\beta^*) := \kappa\left((r+d-x)^2 + r^2 - x^2\right)^{\beta^*/2} - (r+d-x).$$

We consider separately the cases where $\beta^* = 2$ and $\beta^* \in (1,2)$. First consider $\beta^* = 2$. Let $\hat{\eta} > 0$ be such that (41) holds whenever $\max\{r,d\} \leq \hat{\eta}$, and let any $\bar{c} = 1$ and $\eta = \min\{\hat{\eta}, \bar{\eta}/4, \kappa/4\}$, so that $r = \eta$ and $d = 1 - \delta$. For sufficiently large δ, we have $2r + d \leq \bar{\eta}$ and $d \leq \hat{\eta}$, and hence (40) and (41) hold. In addition, since $f(x,\delta,2)$ is linear in x when $\beta^* = 2$, (42) holds whenever $f(r,\delta,2) \le 0$ and $f(-r,\delta,2) \le 0$. In turn, these inequalities hold for sufficiently large δ, since $f(r,\delta,2) = d(\kappa d - 1)$ and $\lim_{\delta\to1} \kappa d - 1 < 0$, and $f(-r,\delta,\beta^*) = (2r+d)(\kappa(2r+d)-1)$ and $\lim_{\delta\to1}\kappa(2r+d)-1 = 2\kappa\eta - 1 < 0$.

Next, consider $\beta^* \in (1,2)$. Let $\bar{c} = 2^{2/(2-\beta^*)}\kappa^{2/(2-\beta^*)}\beta^{*\beta^*/(2-\beta^*)}$ and $\eta = 1$, so that $r = (1-\delta)^{1-\beta^*/2}$ and $d = \bar{c}(1-\delta)^{\beta^*/2}$. Since $\max\{r,d\} \to 0$ as $\delta \to 1$, (40) and (41) hold for sufficiently large δ. In addition, $f(x,\delta,\beta^*)$ is concave in x and is maximized over $x \in [-r,r]$ at

$$x^* = \frac{2r^2 + 2dr + d^2 - (\kappa(r+d)\beta^*)^{\frac{2}{2-\beta^*}}}{2(r+d)}.$$

It thus suffices to show that $f(x^*,\delta,\beta^*) \le 0$ for sufficiently large δ. By algebra,

$$f(x^*,\delta,\beta^*) = -\frac{2r+d}{r+d}\frac{d}{2} + \left(\beta^{*\frac{\beta^*}{2-\beta^*}} - \frac{1}{2}\beta^{*\frac{2}{2-\beta^*}}\right)\kappa^{\frac{2}{2-\beta^*}}(r+d)^{\frac{\beta^*}{2-\beta^*}}.$$

Finally, since $r = (1-\delta)^{1-\beta^*/2} \ge \bar{c}(1-\delta)^{\beta^*/2} = d$ for sufficiently large δ, we have

$$\begin{aligned} f(x^*,\delta,\beta^*) &\le -\frac{d}{2} + 2^{\frac{\beta^*}{2-\beta^*}}\kappa^{\frac{2}{2-\beta^*}}\beta^{*\frac{\beta^*}{2-\beta^*}} r^{\frac{\beta^*}{2-\beta^*}} \\ &= -\frac{\bar{c}(1-\delta)^{\frac{\beta^*}{2}}}{2} + 2^{\frac{\beta^*}{2-\beta^*}}\kappa^{\frac{2}{2-\beta^*}}\beta^{*\frac{\beta^*}{2-\beta^*}}(1-\delta)^{\left(1-\frac{\beta^*}{2}\right)\frac{\beta^*}{2-\beta^*}} \\ &= (1-\delta)^{\frac{\beta^*}{2}}\left(-\frac{\bar{c}}{2} + 2^{\frac{\beta^*}{2-\beta^*}}\kappa^{\frac{2}{2-\beta^*}}\beta^{*\frac{\beta^*}{2-\beta^*}}\right) = 0. \end{aligned}$$

■

We now complete the proof of Lemma 3. Take $\bar{c}$, η, $\bar{\delta}$, B, and r as in Lemma 17. Let B' be the radial contraction of B by a factor of $1 - 72\bar{x}^2(1-\delta)^{\beta^*/2}/(\eta r)$, and define $\varepsilon = 72\bar{x}^2(1-\delta)^{\beta^*/2}/\eta$ and $c = \bar{c} + 72\bar{x}^2/\eta$. Since $d(v,B) = \bar{c}(1-\delta)^{\beta^*/2}$, we have $d(v,B') = (\bar{c} + 72\bar{x}^2/\eta)(1-\delta)^{\beta^*/2} = c(1-\delta)^{\beta^*/2}$, so (39) holds. Moreover, denoting the radius of B' by r', we have

$$r'\varepsilon = \left(1 - \frac{72\bar{x}^2(1-\delta)^{\beta^*/2}}{\eta^2(1-\delta)^{1-\beta^*/2}}\right)\eta(1-\delta)^{1-\beta^*/2} \times \frac{72\bar{x}^2(1-\delta)^{\beta^*/2}}{\eta} = \left(1 - \frac{72\bar{x}^2(1-\delta)^{\beta^*-1}}{\eta^2}\right)72\bar{x}^2(1-\delta).$$

For sufficiently large δ, this is greater than $36\bar{x}^2(1-\delta)$, so (38) holds. Finally, since $B \subseteq \bar{F}$, for all $\lambda' \in \Lambda$ we have $\max_{v'\in\bar{F}} \lambda' \cdot v' \ge \max_{v'\in B} \lambda' \cdot v' = \max_{v'\in B'} \lambda' \cdot v' + 72\bar{x}^2(1-\delta)^{\beta^*/2}/\eta = \max_{v'\in B'} \lambda' \cdot v' + \varepsilon$, so (37) holds.

F Proof of Proposition 2

To define $\bar{x}$, we first observe that for each pair of players $i \neq j$ and each action profile a, we can take $\left(x_i^j(d;a)\right)_d$ such that (i) $\sum_d x_i^j(d;a)\,y^d$ has mean 0 and bounded Euclidean norm; (ii) rewards $\sum_d x_i^j(d;a)\,y^d$ induce player i to take a_i when her opponents take a_{-i}; and (iii) $\mathbb{E}\left[\sum_d x_i^j(d;a)\,y^d|a\right]$ is independent of player j's action.

Lemma 18 *There exists $\hat{x}$ such that, for each pair of players $i \neq j$ and action profile $a \in A$, there exist $\left(x_i^j(d;a)\right)_d$ such that $\mathbb{E}\left[\sum_d x_i^j(d;a)\,y^d|a\right] = 0$, $\frac{d}{da_i}\mathbb{E}\left[\sum_d x_i^j(d;a)\,y^d|a\right] = 1$, $\frac{d}{da_j}\mathbb{E}\left[\sum_d x_i^j(d;a)\,y^d|a\right] = 0$, and $\left|\sum_d x_i^j(d;a)\,y^d\right| \leq \hat{x}$ for all y.*

Proof. For each a and (i,j), let $f^{ij}(a)$ be the value of the program

$$\inf_{b\in\mathbb{R}^D} |b| \qquad \text{subject to}$$

$$\begin{aligned}\sum_d b_d \frac{d}{da_i}\mu(a_i,a_{-i}) &= 1, \text{ or equivalently } bM_i(a) = 1,\\ \sum_d b_d \frac{d}{da_j}\mu(a_i,a_{-i}) &= 0, \text{ or equivalently } bM_j(a) = 0.\end{aligned}$$

(Here b is a row vector while $M_i(a)$ and $M_j(a)$ are column vectors.)

Since $A \ni a$ is compact and N is finite, it suffices to prove that, for each (i,j), (i) $f^{ij}(a) < \infty$ for all a, and (ii) $f^{ij}(a)$ is upper-semicontinuous.

We first prove (i). As in Lemma 1 of Sannikov (2007), pairwise identifiability implies that the columns of $[M^i(a);M^j(a)]$ are linearly independent, so there exists $L(a)$ such that $[M^i(a);M^j(a);L(a)]$ is a D-dimensional invertible matrix. For

$$Q(a) = [M_i(a);0;0]\left[M^i(a);M^j(a);L(a)\right]^{-1},$$

we have $Q(a)M^i(a) = M^i(a)$ and $Q(a)M^j(a) = 0$. Moreover, since $M^i(a)$ is non-degenerate, there exists $\bar{b}$ such that $\bar{b}M^i(a) = 1$. Since $b = \bar{b}Q(a)$ satisfies the constraints, we have $f^{ij}(a) < \infty$.

We next prove (ii). Fix any a and η_0. There exists b such that $|b| \leq f^{ij}(a) + \frac{\eta_0}{2}$ and b satisfies $bM_i(a) = 1$ and $bM_j(a) = 0$. Take $L(a)$ as in the proof of (i). Taking $\eta_1 > 0$ sufficiently small, we can guarantee that $[M^i(a');M^j(a');L(a)]$ is a D-dimensional invertible matrix for each a' with $|a - a'| \leq \eta_1$. Define a D-dimensional vector $\Delta_{a'}$ by

$$\Delta_{a'} = [b(M_i(a') - M_i(a)), b(M_j(a') - M_j(a)), 0]\left[M^i(a');M^j(a');L(a)\right]^{-1}.$$

By definition,

$$
\begin{aligned}
\left(b+\Delta_{a'}\right) M_i\left(a'\right) &= bM_i\left(a'\right)-b\left(M_i\left(a'\right)-M_i\left(a\right)\right)=bM_i\left(a\right)=1, \\
\left(b+\Delta_{a'}\right) M_j\left(a'\right) &= bM_j\left(a'\right)-b\left(M_j\left(a'\right)-M_j\left(a\right)\right)=bM_j\left(a\right)=0.
\end{aligned}
$$

Thus, $b-\Delta_{a'}$ satisfies the constraint for a', and hence $f^{ij}\left(a'\right) \leq |b|+\left|\Delta_{a'}\right|$. Since $\limsup_{\eta_1 \to 0} \sup_{a':|a-a'| \leq \eta_1} \left|\Delta_{a'}\right|=0$, for sufficiently small $\eta_1>0$, we have $f^{ij}\left(a'\right) \leq |b|+\left|\Delta_{a'}\right| \leq |b|+\frac{1}{2}\eta_0 \leq f^{ij}\left(a\right)+\eta_0$ for all a' with $\left|a-a'\right| \leq \eta_1$, establishing upper-semicontinuity. ■

Given Lemma 18, Assumption 2(ii) holds with $\bar{x}=\bar{u}\hat{x}$. To see why, for any i and a, let $\partial u_i=\frac{\partial}{\partial a_i'} u_i\left(a_i', a_{-i}\right)\Big|_{a_i'=a_i}$ and $x_i\left(y\right)=-\partial u_i \sum_d x_i^j\left(d; a^\lambda\right) y^d$. Then,

$$
\begin{aligned}
\frac{\partial}{\partial a_i'}\left(u_i\left(a_i', a_{-i}\right)+\mathbb{E}\left[x_i\left(y\right)|a_i', a_{-i}\right]\right)\Big|_{a_i'=a_i} &= 0 \qquad \text{and} \\
\frac{\partial}{\partial a_j'}\mathbb{E}\left[cx_i\left(y\right)|a_j', a_{-i}\right]\Big|_{a_j'=a_j} &= 0 \qquad \text{for all } j \neq i.
\end{aligned}
$$

Since u_i is concave in a_i, $\mathbb{E}\left[x_i\left(y\right)|a_i', a_{-i}\right]$ is linear in a_i', and $\mathbb{E}\left[cx_i\left(y\right)|a_j', a_{-j}\right]$ is linear in a_j', we have (6) and (7). Moreover, since $\left|\partial u_i\right| \leq \bar{u}$, we have $\left|x_i\left(y\right)\right| \leq \bar{u}\hat{x}$ for all i, y.

G Omitted Details for the Proof of Theorem 3

We require some preliminary lemmas. The first two derive properties of the feasible payoff set.

Lemma 19 *There exists $\varepsilon_1>0$ such that, for any $\alpha \in \Delta\left(A\right)$, we have*

$$
\mathbb{E}^\alpha\left[a\right]-u^{-1}\left(w+\mathbb{E}^\alpha\left[\psi\left(a\right)\right]\right) \leq \bar{F}\left(w\right)-\varepsilon_1 \mathbb{E}^\alpha\left[\left(a-\bar{a}\left(w\right)\right)^2\right].
$$

Proof. Since $\psi \in C^2$ and $\psi''>0$, there exists $\varepsilon_1>0$ such that, for all $a \in A$, we have

$$
\psi\left(a\right)-\psi\left(\bar{a}\left(w\right)\right) \geq \psi'\left(\bar{a}\left(w\right)\right)\left(a-\bar{a}\left(w\right)\right)+\varepsilon_1\left(a-\bar{a}\left(w\right)\right)^2.
$$

Thus, for any $\alpha \in \Delta(A)$, we have

$$\begin{aligned}
& \bar{F}(w) - \mathbb{E}^{\alpha}[a] + u^{-1}(w + \mathbb{E}^{\alpha}[\psi(a)]) \\
= \;& \bar{a}(w) - u^{-1}(w + \psi(\bar{a}(w))) - \mathbb{E}^{\alpha}[a] + u^{-1}(w + \mathbb{E}^{\alpha}[\psi(a)]) \\
\geq \;& \bar{a}(w) - \mathbb{E}^{\alpha}[a] + \frac{\mathbb{E}^{\alpha}[\psi(a)] - \psi(\bar{a}(w))}{u'(\bar{c}(w))} \\
\geq \;& \bar{a}(w) - \mathbb{E}^{\alpha}[a] + \frac{\psi'(\bar{a}(w))(\mathbb{E}^{\alpha}[a] - \bar{a}(w)) + \varepsilon_1 \mathbb{E}^{\alpha}\left[(a - \bar{a}(w))^2\right]}{u'(\bar{c}(w))} \geq \frac{\varepsilon_1 \mathbb{E}^{\alpha}\left[(a - \bar{a}(w))^2\right]}{u'(\bar{c}(w))},
\end{aligned}$$

where the first inequality is by Taylor expansion of $u^{-1}(w + \mathbb{E}^{\alpha}[\psi(a)]) - u^{-1}(w + \psi(\bar{a}(w)))$ around $w + \mathbb{E}^{\alpha}[\psi(a)] = w + \psi(\bar{a}(w)) = \bar{c}(w)$, the second inequality is by Taylor expansion of $\psi(a) - \psi(\bar{a}(w))$ around $a = \bar{a}(w)$, and the last equality is by $u'(\bar{c}(w)) = \psi'(\bar{a}(w))$ (by definition of $\bar{a}(w)$). ■

Lemma 20 *There exists $\varepsilon_2 > 0$ such that $\bar{F}''(w) \leq -\varepsilon_2$ for all $w \geq -\psi(\bar{A})$.*

Proof. Differentiating the equality $u'(\bar{c}(w)) = \psi'(\bar{a}(w))$ with respect to w yields

$$\bar{a}'(w) = \frac{u''(\bar{c}(w))}{u'(\bar{a}(w))(\psi''(\bar{a}(w)) - u''(\bar{c}(w)))}. \tag{43}$$

Since $\bar{F}(w) = \max_{a \in A} a - u^{-1}(w + \psi(a))$, by the envelope theorem we have $\bar{F}'(w) = -1/u'(\bar{c}(w))$, or equivalently $\bar{F}'(w) = -1/u'(w + \psi(\bar{a}(w)))$. Differentiating this equality respect to w and substituting (43) yields

$$\bar{F}''(w) = \frac{u''(\bar{c}(w))\psi''(\bar{a}(w))}{u'(\bar{c}(w))^3(\psi''(\bar{a}(w)) - u''(\bar{c}(w)))}.$$

The lemma follows since $\psi'' > 0$, $u'' < 0$, and $u''(c)/u'(c)^3$ is bounded away from zero by (10). ■

The next lemma gives a key probability bound.

Lemma 21 *For any $c > 0$ and $\vartheta \geq 0$, there exists $\bar{\delta} < 1$ such that, for any $\delta > \bar{\delta}$ and any sequence of non-negative random variables $(X_t)_{t \geq 1}$, where X_t is distributed according to a cdf G_t satisfying $1 - G_t(x) \leq 2\exp\left(-cx^2/\left((1-\delta)^2 t\right)\right)$ for all t, we have*

$$(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}t\int_{x \geq (1-\delta)^{\frac{1}{2}-\varepsilon}} x^{\vartheta} dG_t(x) \leq (1-\delta)^2.$$

Proof. Let $T = (1-\delta)^{\frac{1}{2}-\varepsilon}$. It suffices to show that

$$\lim_{\delta\to1}(1-\delta)^{-1}\sum_{t=1}^{T}\delta^{t-1}t\int_{x\geq(1-\delta)^{\frac{1}{2}-\varepsilon}}x^{\vartheta}dG_t(x)=\lim_{\delta\to1}(1-\delta)^{-1}\sum_{t=T+1}^{\infty}\delta^{t-1}t\int_{x\geq(1-\delta)^{\frac{1}{2}-\varepsilon}}x^{\vartheta}dG_t(x)=0.$$

For each t, we have

$$\begin{aligned}&\int_{x\geq(1-\delta)^{\frac{1}{2}-\varepsilon}}x^{\vartheta}dG_t(x)\\ =\ &\int_{x\geq(1-\delta)^{\frac{1}{2}-\varepsilon}}\vartheta x^{\vartheta-1}(1-G_t(x))\,dx+(1-\delta)^{\vartheta\left(\frac{1}{2}-\varepsilon\right)}\left(1-G_t\left((1-\delta)^{\frac{1}{2}-\varepsilon}\right)\right)\\ \leq\ &2\vartheta\int_{x\geq(1-\delta)^{\frac{1}{2}-\varepsilon}}x^{\vartheta-1}\exp\left(-\frac{cx^2}{(1-\delta)^2t}\right)dx+2(1-\delta)^{\vartheta\left(\frac{1}{2}-\varepsilon\right)}\exp\left(-\frac{c(1-\delta)^{1-2\varepsilon}}{(1-\delta)^2t}\right),\end{aligned}$$

where the equality is by integration by parts and $\lim_{x\to\infty}x^{\vartheta}(1-G_t(x))=0$, and the inequality is by $1-G_t(x)\leq2\exp\left(-cx^2/\left((1-\delta)^2t\right)\right)$. Note that $\sum_{t=1}^{T}t\delta^{t-1}\leq(1-\delta)^{-2}$ and, if $t\leq T$, then $(1-\delta)^{1-2\varepsilon}/\left((1-\delta)^2t\right)\geq(1-\delta)^{-\varepsilon}$. Therefore

$$\lim_{\delta\to1}(1-\delta)^{-1}\sum_{t=1}^{T}t\delta^{t-1}(1-\delta)^{\vartheta\left(\frac{1}{2}-\varepsilon\right)}\exp\left(-\frac{c(1-\delta)^{1-2\varepsilon}}{(1-\delta)^2t}\right)=0.$$

At the same time, since $\sum_{t=T+1}^{\infty}t\delta^{t-1}=\frac{\delta^{T+1}(1+T(1-\delta))}{\delta(1-\delta)^2}\leq(1-\delta)^{-2-\varepsilon}\delta^{T}\leq(1-\delta)^{-2-\varepsilon}\exp\left(-(1-\delta)^{-\varepsilon}\right)$, we have

$$\lim_{\delta\to1}(1-\delta)^{-1}\sum_{t=T+1}^{\infty}t\delta^{t-1}(1-\delta)^{\vartheta\left(\frac{1}{2}-\varepsilon\right)}\exp\left(-\frac{c(1-\delta)^{1-2\varepsilon}}{(1-\delta)^2t}\right)=0.$$

It thus suffices to show that

$$\lim_{\delta\to1}(1-\delta)^{-1}\sum_{t=1}^{T}t\delta^{t-1}\int_{x\geq(1-\delta)^{\frac{1}{2}-\varepsilon}}x^{\vartheta-1}\exp\left(-\frac{cx^2}{(1-\delta)^2t}\right)dx=0,\quad\text{and}\tag{44}$$

$$\lim_{\delta\to1}(1-\delta)^{-1}\sum_{t=T+1}^{\infty}t\delta^{t-1}\int_{x\geq(1-\delta)^{\frac{1}{2}-\varepsilon}}x^{\vartheta-1}\exp\left(-\frac{cx^2}{(1-\delta)^2t}\right)dx=0.\tag{45}$$

We first establish (44). Since $\sum_{t=1}^{T}t\delta^{t-1}\leq(1-\delta)^{-2}$ and $(1-\delta)^2T=(1-\delta)^{1-\varepsilon}$, it suffices to show that

$$\lim_{\delta\to1}(1-\delta)^{-3}\int_{x\geq(1-\delta)^{\frac{1}{2}-\varepsilon}}x^{\vartheta-1}\exp\left(-\frac{cx^2}{(1-\delta)^{1-\varepsilon}}\right)dx=0.$$

If $\vartheta \leq 1$ then

$$\begin{aligned}
& (1-\delta)^{-3} \int_{x \geq (1-\delta)^{\frac{1}{2}-\varepsilon}} x^{\vartheta-1} \exp\left(-\frac{cx^2}{(1-\delta)^{1-\varepsilon}}\right) dx \\
\leq \; & (1-\delta)^{\left(\frac{1}{2}-\varepsilon\right)(\vartheta-1)-3} \int_{x \geq (1-\delta)^{\frac{1}{2}-\varepsilon}} \exp\left(-\frac{cx^2}{(1-\delta)^{1-\varepsilon}}\right) dx \\
\leq \; & (1-\delta)^{\left(\frac{1}{2}-\varepsilon\right)(\vartheta-1)-3} \frac{\sqrt{\pi}}{c} (1-\delta)^{1/2} \exp\left(-c(1-\delta)^{-\varepsilon}\right) \xrightarrow{\delta \to 1} 0,
\end{aligned}$$

where the second inequality follows by the same calculation as (18). If instead $\vartheta > 1$ then, for sufficiently large δ, we have

$$\begin{aligned}
& (1-\delta)^{-3} \int_{x \geq (1-\delta)^{\frac{1}{2}-\varepsilon}} x^{\vartheta-1} \exp\left(-\frac{cx^2}{(1-\delta)^{1-\varepsilon}}\right) dx \\
= \; & \frac{1}{2}(1-\delta)^{-3} \int_{y \geq 0} \left(y + (1-\delta)^{1-2\varepsilon}\right)^{\vartheta-1} \exp\left(-\frac{c\left(y + (1-\delta)^{1-2\varepsilon}\right)}{(1-\delta)^{1-\varepsilon}}\right) dy \\
\leq \; & \frac{1}{2}(1-\delta)^{-3+\vartheta(1-2\varepsilon)} \exp\left(-c(1-\delta)^{-\varepsilon}\right) \int_{y \geq 0} \exp\left(\left(\frac{\vartheta-1}{(1-\delta)^{1-2\varepsilon}} - \frac{c}{(1-\delta)^{1-\varepsilon}}\right) y\right) dy \\
= \; & \frac{1}{2}(1-\delta)^{-3+\vartheta(1-2\varepsilon)} \exp\left(-c(1-\delta)^{-\varepsilon}\right) \left(-\frac{\vartheta-1}{(1-\delta)^{1-2\varepsilon}} + \frac{c}{(1-\delta)^{1-\varepsilon}}\right)^{-1} \xrightarrow{\delta \to 1} 0,
\end{aligned}$$

where the second follows by integration by substitution (setting $y = x^2 - (1-\delta)^{1-2\varepsilon}$), the third line follows because

$$\begin{aligned}
\left(y + (1-\delta)^{1-2\varepsilon}\right)^{\vartheta-1} &= (1-\delta)^{(1-2\varepsilon)(\vartheta-1)} \exp\left((\vartheta-1)\log\left(\frac{y}{(1-\delta)^{1-2\varepsilon}} + 1\right)\right) \\
&\leq (1-\delta)^{(1-2\varepsilon)(\vartheta-1)} \exp\left(\frac{(\vartheta-1)\, y}{(1-\delta)^{1-2\varepsilon}}\right),
\end{aligned}$$

and the fourth line follows because $\frac{\vartheta-1}{(1-\delta)^{1-2\varepsilon}} - \frac{c}{(1-\delta)^{1-\varepsilon}} < 0$ for sufficiently large δ.

We next establish (45). For any t, we have

$$\begin{aligned}
\int_{x \geq 0} x^{\vartheta-1} \exp\left(-\frac{cx^2}{(1-\delta)^2 t}\right) dx &= \frac{1}{2}\left(\frac{(1-\delta)^2 t}{c}\right)^{\frac{\vartheta}{2}} \int_{y \geq 0} y^{\frac{\vartheta}{2}-1} \exp(-y)\, dy \\
&= \frac{1}{2}\left(\frac{(1-\delta)^2 t}{c}\right)^{\frac{\vartheta}{2}} \Gamma\left(\frac{\vartheta}{2}\right),
\end{aligned}$$

where the first line follows by setting $y = cx^2 / \left((1-\delta)^2 t \right)$, and the second line follows by the definition of the gamma function, Γ. Hence, there exist constants c_1, c_2 such that

$$\int_{x \geq 0} x^{\vartheta - 1} \exp \left(-\frac{cx^2}{(1-\delta)^2 t} \right) dx \leq c_1 (1-\delta)^{2c_2} t^{c_2} \qquad \text{for all } t.$$

We thus have

$$\begin{aligned}
& (1-\delta)^{-1} \sum_{t=T+1}^{\infty} t \delta^{t-1} \int_{x \geq (1-\delta)^{\frac{1}{2}-\varepsilon}} x^{\vartheta-1} \exp \left(-\frac{cx^2}{(1-\delta)^2 t} \right) dx \\
\leq \; & (1-\delta)^{-1} \sum_{t=T+1}^{\infty} t \delta^{t-1} \int_{x \geq 0} x^{\vartheta-1} \exp \left(-\frac{cx^2}{(1-\delta)^2 t} \right) dx \\
\leq \; & c_1 (1-\delta)^{2c_2 - 1} \sum_{t=T+1}^{\infty} \delta^{t-1} t^{c_2+1} \\
\leq \; & c_1 (1-\delta)^{2c_2-1} \delta^{T-1} (T+1)^{c_2+1} \sum_{t=1}^{\infty} \delta^t t^{c_2} \\
\leq \; & c_1 (1-\delta)^{c_2-2} \delta^{T-1} (T+1)^{c_2+1} k(c_2) \overset{\delta \to 1}{\longrightarrow} 0,
\end{aligned}$$

where k is defined in (20) and the limit follows recalling that $\delta^T \leq \exp \left(-(1-\delta)^{-\varepsilon} \right)$. ■

We now establish inequality (15).

Lemma 22 *We have*

$$\mathbb{E}^{\alpha_t} \left[\frac{\psi'(a)^2}{\mathcal{I}(a)} \right] \leq \mathbb{E}^{\mu} \left[\frac{\psi'(a_t) \nu(y_t | a_t)}{\mathcal{I}(a_t)} \sum_{t'=t+1}^{\infty} \delta^{t'-t} \hat{u}_{t'} \right] \qquad \textit{for all } t.$$

Proof. Fix any t and $\varepsilon > 0$. Since $\psi'(a) / \mathcal{I}(a)$ is bounded (uniformly in a) given Assumption 3(ii), there exists $\tilde{\Delta} > 0$ such that, for all $\Delta < \tilde{\Delta}$, we have $(\psi'(a) / \mathcal{I}(a)) \Delta \leq \varepsilon$ for all $a \in A$. Fix $\bar{\Delta} < \tilde{\Delta}$ such that, for all $a \in A$, $\int_y \left(\max_{\tilde{a} \in [a, a+\bar{\Delta}]} p_a(y|\tilde{a})^2 / p(y|a) \right) dy < \infty$. (Such $\bar{\Delta}$ exists by Assumption 3(i).) Now, for any $\Delta < \bar{\Delta}$, consider the manipulation where, whenever the agent is recommended action $a \geq \varepsilon$ in period t, she instead takes $a - (\psi'(a) / \mathcal{I}(a)) \Delta$. This manipulation is unprofitable for the agent if and only if

$$\begin{aligned}
& \mathbb{E}^{\alpha_t} \left[\mathbf{1}\{a \geq \varepsilon\} \left(\psi(a) - \psi \left(a - \frac{\psi'(a)}{\mathcal{I}(a)} \Delta \right) \right) \right] \\
& \qquad \leq \mathbb{E}^{\mu} \left[\mathbf{1}\{a_t \geq \varepsilon\} \frac{p(y_t|a_t) - p\left(y_t | a_t - \frac{\psi'(a_t)}{\mathcal{I}(a_t)} \Delta \right)}{p(y|a_t)} \sum_{t'=t+1}^{\infty} \delta^{t'-t} \hat{u}_{t'} \right].
\end{aligned}$$

Since this holds for all $\Delta < \bar{\Delta}$, we have

$$\lim_{\Delta \to 0} \mathbb{E}^{\alpha_t} \left[\mathbf{1}\left\{a \geq \varepsilon\right\} \frac{\psi\left(a\right) - \psi\left(a - \frac{\psi'(a)}{\mathcal{I}(a)}\Delta\right)}{\Delta} \right]$$
$$\leq \lim_{\Delta \to 0} \mathbb{E}^{\mu} \left[\mathbf{1}\left\{a_t \geq \varepsilon\right\} \frac{p\left(y_t|a_t\right) - p\left(y_t|a_t - \frac{\psi'(a_t)}{\mathcal{I}(a_t)}\Delta\right)}{\Delta p\left(y_t|a_t\right)} \sum_{t'=t+1}^{\infty} \delta^{t'-t} \hat{u}_{t'} \right]$$

In this inequality, the LHS is bounded because $\psi'\left(a\right)^2 / \mathcal{I}\left(a\right)$ is bounded given Assumption 3(ii), and hence is equal to $\mathbb{E}^{\alpha_t}\left[\mathbf{1}\left\{a \geq \varepsilon\right\} \psi'\left(a\right)^2 / \mathcal{I}\left(a\right)\right]$, by dominated convergence. As for the RHS, for each $t' \geq t+1$, by Cauchy-Schwarz the corresponding term in the sum is bounded by

$$\mathbb{E}^{\mu} \left[\mathbf{1}\left\{a_t \geq \varepsilon\right\} \left(\frac{p\left(y_t|a_t\right) - p\left(y_t|a_t - \frac{\psi'(a_t)}{\mathcal{I}(a_t)}\Delta\right)}{\Delta p\left(y|a_t\right)} \right)^2 \right]^{1/2} \times \mathbb{E}^{\mu}\left[\hat{u}_{t'}^2\right]^{1/2}.$$

This is finite because $\int_y \left(\max_{\tilde{a} \in \left[a, a+\bar{\Delta}\right]} p_a\left(y|\tilde{a}\right)^2 / p\left(y|a\right)\right) dy < \infty$ (by Assumption 3(i)) and $\mathbb{E}^{\mu}\left[\hat{u}_{t'}^2\right] < \infty$ for all t' (as otherwise the principal's expected payoff would equal $-\infty$ by Lemma 20). Hence, the entire RHS is bounded, and hence is equal to $\mathbb{E}^{\mu}\Big[\mathbf{1}\left\{a_t \geq \varepsilon\right\} \left(\psi'\left(a_t\right) \nu\left(y_t|a_t\right) / \mathcal{I}\left(a_t\right)\right)$ by dominated convergence. In total, we have

$$\mathbb{E}^{\alpha_t} \left[\mathbf{1}\left\{a \geq \varepsilon\right\} \frac{\psi'\left(a\right)^2}{\mathcal{I}\left(a\right)} \right] \leq \mathbb{E}^{\mu} \left[\mathbf{1}\left\{a_t \geq \varepsilon\right\} \frac{\psi'\left(a_t\right) \nu\left(y_t|a_t\right)}{\mathcal{I}\left(a_t\right)} \sum_{t'=t+1}^{\infty} \delta^{t'-t} \hat{u}_{t'} \right].$$

Since this holds for all $\varepsilon > 0$, and $\psi'\left(a\right)^2 / \mathcal{I}\left(a\right)$ and $\psi'\left(a_t\right) \nu\left(y_t|a_t\right) / \mathcal{I}\left(a_t\right)$ are continuous, taking $\varepsilon \to 0$ completes the proof. ■

Now we prove our key lemmas, Lemmas 10 and 11. These complete the proof that inefficiency is at least (14) in the blind game.

G.1 Proof of Lemma 10

Multiplying both sides of (28) by $2\delta/((1-\delta)\bar{F}''\left(w\right))$ and using (27), it suffices to find $c, \varepsilon > 0$ such that

$$\min_{\left(\hat{u}_t(h)\right)_{t,h} \geq -\left(w + \psi\left(\bar{A}\right)\right) \text{ s.t. } (15)} \sum_{t=1}^{\infty} \delta^t \frac{2\mathbb{E}^{\mu}\left[\tilde{F}\left(\hat{u}_t\left(h\right)\right)\right]}{\bar{F}''\left(w\right)} \geq \mathbb{E}^{\alpha}\left[\frac{\psi'\left(a\right)^2}{\mathcal{I}\left(a\right)}\right] - c\left(1-\delta\right)^{\varepsilon},$$

where

$$\tilde{F}(\hat{u}) = \bar{F}(w+\hat{u}) - \bar{F}(w) - \bar{F}'(w)\hat{u}.$$

By weak duality, taking a Lagrange multiplier of $2(1-\delta)\delta^{t-1}$ on (15) for each t, the LHS is no less than

$$\min_{(\hat{u}_t(h))_{t,h} \geq -\left(w+\psi\left(\bar{A}\right)\right)} 2\sum_{t=1}^{\infty}\delta^t \mathbb{E}^{\mu}\left[\frac{\tilde{F}(\hat{u}_t(h))}{\bar{F}''(w)} - \underbrace{\left(\frac{1-\delta}{\delta}\sum_{t'=1}^{t-1}\frac{\psi'(a_{t'})\nu(y_t|a_{t'})}{\mathcal{I}(a_{t'})}\right)}_{:=\Omega_t^\delta(h)}\hat{u}_t(h)\right] + 2\mathbb{E}^{\alpha}\left[\frac{\psi'(a)^2}{\mathcal{I}(a)}\right]. \tag{46}$$

It remains to bound (46). Since $\tilde{F}$ is concave, the first-order necessary and sufficient condition for $\hat{u}_t(h)$ is

$$\hat{u}_t(h) = \begin{cases} \tilde{F}'^{-1}\left(\bar{F}''(w)\Omega_t^\delta(h)\right) & \text{if } \tilde{F}'^{-1}\left(\bar{F}''(w)\Omega_t^\delta(h)\right) \geq -\left(w+\psi\left(\bar{A}\right)\right), \\ -\left(w+\psi\left(\bar{A}\right)\right) & \text{if } \tilde{F}'^{-1}\left(\bar{F}''(w)\Omega_t^\delta(h)\right) < -\left(w+\psi\left(\bar{A}\right)\right). \end{cases}$$

Now fix any $\varepsilon \in (0, 1/8)$ and let

$$\bar{H}^t = \left\{h \in (A\times Y)^\infty : -(1-\delta)^{\frac{1}{2}-\varepsilon} \leq \Omega_t^\delta(h) \leq (1-\delta)^{\frac{1}{2}-\varepsilon}\right\} \qquad \text{for all } t. \tag{47}$$

We establish two further lemmas.

Lemma 23 *There exists $\hat{c} > 0$ such that, for any sufficiently large $\delta < 1$, we have*

$$2\left(\frac{\tilde{F}(\hat{u}_t(h))}{\bar{F}''(w)} - \Omega_t^\delta(h)\hat{u}_t(h)\right) \geq -\Omega_t^\delta(h)^2 - \hat{c}(1-\delta)^{1+\varepsilon} \qquad \textit{for all } t \textit{ and } h \in \bar{H}_t. \tag{48}$$

Proof. For sufficiently large δ, we have $\hat{u}_t(h) = \tilde{F}'^{-1}\left(\bar{F}''(w)\Omega_t^\delta(h)\right)$ for all t and $h \in \bar{H}^t$. Since $\tilde{F} \in C^2$ and $\Omega_t^\delta(h)$ is bounded (uniformly in δ) for $h \in \bar{H}^t$, by Taylor expansion of $\tilde{F}'^{-1}$ and $\tilde{F} \circ \tilde{F}'^{-1}$ around 0, there exists $\hat{c} > 0$ such that, for any $\delta < 1$ and $h \in \bar{H}^t$, we have

$$\begin{aligned} \left|\hat{u}_t(h) - \left(\tilde{F}'^{-1}(0) - \frac{\bar{F}''(w)}{\tilde{F}''\circ\tilde{F}'^{-1}(0)}\Omega_t^\delta(h)\right)\right| &\leq \frac{\hat{c}}{3}\Omega_t^\delta(h)^2 \qquad \text{and} \\ \left|\frac{2\tilde{F}(\hat{u}_t(h))}{\bar{F}''(w)} - \left(\begin{array}{c} \frac{2\tilde{F}\circ\tilde{F}'^{-1}(0)}{\bar{F}''(w)} - \frac{2\tilde{F}'\circ\tilde{F}'^{-1}(0)}{\tilde{F}''\circ\tilde{F}'^{-1}(0)}\Omega_t^\delta(h) \\ +\frac{\bar{F}''(w)}{\tilde{F}''\circ\tilde{F}'^{-1}(0)}\left(1 - \frac{\tilde{F}'\circ\tilde{F}'^{-1}(0)\times\tilde{F}'''\circ\tilde{F}'^{-1}(0)}{\left(\tilde{F}''\circ\tilde{F}'^{-1}(0)\right)^2}\right)\Omega_t^\delta(h)^2 \end{array}\right)\right| &\leq \frac{\hat{c}}{3}\left|\Omega_t^\delta(h)\right|^3. \end{aligned}$$

Since $\tilde{F}'^{-1}(0) = \tilde{F}'\circ\tilde{F}'^{-1}(0) = 0$ and $\tilde{F}''\circ\tilde{F}'^{-1}(0) = \bar{F}''(w)$ by definition of $\tilde{F}$, these inequalities simplify to

$$\begin{aligned} \left|\hat{u}_t(h) - \Omega_t^\delta(h)\right| &\le \frac{\hat{c}}{3}\Omega_t^\delta(h)^2 \qquad \text{and} \\ \left|\frac{2\tilde{F}(\hat{u}_t(h))}{\bar{F}''(w)} - \Omega_t^\delta(h)^2\right| &\le \frac{\hat{c}}{3}\left|\Omega_t^\delta(h)\right|^3. \end{aligned}$$

Multiplying the first inequality by $\Omega_t^\delta(h)$ and applying the triangle inequality gives

$$\begin{aligned} \left|\Omega_t^\delta(h)\hat{u}_t(h) - \Omega_t^\delta(h)^2\right| &\le \frac{\hat{c}}{3}\Omega_t^\delta(h)^3 \qquad \text{and} \\ \left|\frac{2\tilde{F}(\hat{u}_t(h))}{\bar{F}''(w)} - \Omega_t^\delta(h)\hat{u}_t(h)\right| &\le \frac{2\hat{c}}{3}\left|\Omega_t^\delta(h)\right|^3. \end{aligned}$$

We thus have

$$\begin{aligned} 2\left(\frac{\tilde{F}(\hat{u}_t(h))}{\bar{F}''(w)} - \Omega_t^\delta(h)\hat{u}_t(h)\right) &\ge -\Omega_t^\delta(h)\hat{u}_t(h) - \frac{2\hat{c}}{3}\left|\Omega_t^\delta(h)\right|^3 \\ &\ge -\Omega_t^\delta(h)^2 - \hat{c}\left|\Omega_t^\delta(h)\right|^3 \\ &\ge -\Omega_t^\delta(h)^2 - \hat{c}(1-\delta)^{\frac{3}{2}-3\varepsilon} \\ &\ge -\Omega_t^\delta(h)^2 - \hat{c}(1-\delta)^{1+\varepsilon}, \end{aligned}$$

where the third line follows by (47), and the fourth line follows because $\varepsilon < 1/8$. ■

Lemma 24 *For any sufficiently large $\delta < 1$, we have*

$$2\sum_{t=1}^{\infty}\delta^t\int_{h\notin\bar{H}^t}\left(\frac{\tilde{F}(\hat{u}_t(h))}{\bar{F}''(w)} - \Omega_t^\delta(h)\hat{u}_t(h)\right)d\mu(h) \ge -(1-\delta)^\varepsilon. \tag{49}$$

Proof. We first show that, for sufficiently large δ,

$$\sum_{t=1}^{\infty}\delta^t\int_{h\notin\bar{H}^t}\Omega_t^\delta(h)^2\,d\mu(h) \le (1-\delta)^{2\varepsilon}. \tag{50}$$

To see this, note that $(\psi'(a_t)\nu(y_t|a_t)/\mathcal{I}(a_t))_t$ is a sequence of martingale increments where, for all $\theta > 0$,

$$\begin{aligned}\mathbb{E}^{\mu}\left[\exp\left(\theta\frac{1-\delta}{\delta}\frac{\psi'(a_t)\,\nu(y_t|a_t)}{\mathcal{I}(a_t)}\right)|h^t\right] &\le \max_{a\in A}\mathbb{E}^{y\sim p(y|a)}\left[\exp\left(\theta\frac{1-\delta}{\delta}\frac{\psi'(a)\,\nu(y|a)}{\mathcal{I}(a)}\right)\right]\\ &\le \exp\left(\frac{\theta^2}{2}\left(\frac{1-\delta}{\delta}\right)^2\max_{a\in A}\frac{\psi'(a)^2}{\mathcal{I}(a)}\right),\end{aligned}$$

where the second line follows from Assumption 3(iii), noting that $\psi'(a)^2/\mathcal{I}(a)$ is bounded. Hence, by Lemma 9, there exists $\tilde{c} > 0$ such that, for all t and $x \ge 0$, we have

$$\Pr^{\mu}\left(\left|\Omega_t^{\delta}(h)\right| \ge x\right) = \Pr^{\mu}\left(\left|\frac{1-\delta}{\delta}\sum_{t'=1}^{t-1}\frac{\psi'(a_{t'})\,\nu(y_{t'}|a_{t'})}{\mathcal{I}(a_{t'})}\right| \ge x\right) \le 2\exp\left(-\frac{\tilde{c}x^2}{(1-\delta)^2 t}\right).$$

We can thus apply Lemma 21 to the sequence $\left(\left|\Omega^{\delta}\right|_t\right)_{t\ge 1}$ to conclude that (50) holds for sufficiently large δ.

Thus, for sufficiently large δ, we have

$$\begin{aligned}&2\sum_{t=1}^{\infty}\delta^t\int_{h\notin\bar{H}^t}\left(\frac{\tilde{F}(\hat{u}_t(h))}{\bar{F}''(w)} - \Omega_t^{\delta}(h)\,\hat{u}_t(h)\right)d\mu(h)\\ \ge\ &2\sum_{t=1}^{\infty}\delta^t\int_{h\notin\bar{H}^t}\left(-\frac{\varepsilon_2\hat{u}_t(h)^2}{2\bar{F}''(w)} - \Omega_t^{\delta}(h)\,\hat{u}_t(h)\right)d\mu(h)\\ \ge\ &\frac{\bar{F}''(w)}{\varepsilon_2}\sum_{t=1}^{\infty}\delta^t\int_{h\notin\bar{H}^t}\Omega_t^{\delta}(h)^2\,d\mu(h) \ge \frac{\bar{F}''(w)}{\varepsilon_2}(1-\delta)^{2\varepsilon} \ge -(1-\delta)^{\varepsilon},\end{aligned}$$

where the first inequality follows by Lemma 20 (as taking a second-order Taylor expansion gives $\tilde{F}(x) \le -(\varepsilon_2/2)\,x^2$ for all $x \ge -w$), the second follows by minimizing over $x_t(h)$, the third follows by (50), and the fourth follows by $\bar{F}'' < 0$. ■

By (48), (49), and $\Omega_t^{\delta}(h)^2 > 0$, we see that (46) is no less than

$$-\sum_{t=2}^{\infty}\delta^t\mathbb{E}^{\mu}\left[\Omega_t^{\delta}(h)^2\right] + 2\mathbb{E}^{\alpha}\left[\frac{\psi'(a)^2}{\mathcal{I}(a)}\right] - \max\{\hat{c},1\}(1-\delta)^{\varepsilon}.$$

Since $\left(\psi'(a_t)\,\nu(y_t|a_t)/\mathcal{I}(a_t)\right)_t$ is a sequence of martingale increments, we have

$$\mathbb{E}_{\mu}\left[\Omega_t^{\delta}(h)^2\right] = \left(\frac{1-\delta}{\delta}\right)^2\sum_{t'=1}^{t-1}\mathbb{E}^{\mu}\left[\left(\frac{\psi'(a_{t'})\,\nu(y_{t'}|a_{t'})}{\mathcal{I}(a_{t'})}\right)^2\right] = \left(\frac{1-\delta}{\delta}\right)^2\sum_{t'=1}^{t-1}\mathbb{E}^{\alpha_{t'}}\left[\frac{\psi'(a)^2}{\mathcal{I}(a)}\right].$$

Therefore, (46) is no less than

$$
\begin{aligned}
& -\left(\frac{1-\delta}{\delta}\right)^2 \sum_{t=2}^{\infty} \delta^t \sum_{t'=1}^{t-1} \mathbb{E}^{\alpha_{t'}} \left[\frac{\psi'(a)^2}{\mathcal{I}(a)}\right] + 2\mathbb{E}^{\alpha} \left[\frac{\psi'(a)^2}{\mathcal{I}(a)}\right] - \max\{\hat{c}, 1\}(1-\delta)^{\varepsilon} \\
= \; & \mathbb{E}^{\alpha} \left[\frac{\psi'(a)^2}{\mathcal{I}(a)}\right] - \max\{\hat{c}, 1\}(1-\delta)^{\varepsilon}.
\end{aligned}
$$

Taking $c = \max\{\hat{c}, 1\}$ completes the proof.

G.2 Proof of Lemma 11

If α assigns probability 1 to $a = \bar{a}(w)$ then

$$
\begin{aligned}
& \min\left\{\frac{1-\delta}{\delta}\mathbb{E}^{\alpha}\left[\frac{\psi'(a)^2}{\mathcal{I}(a)}\right]\frac{\bar{F}''(w)}{2} + c(1-\delta)^{1+\varepsilon}, -\varepsilon_1 \mathbb{E}^{\alpha}\left[(a-\bar{a}(w))^2\right]\right\} \\
\geq \; & \frac{1-\delta}{\delta}\frac{\psi'(\bar{a}(w))^2}{\mathcal{I}(\bar{a}(w))}\frac{\bar{F}''(w)}{2}.
\end{aligned}
$$

Hence, the optimal α satisfies

$$
\mathbb{E}^{\alpha}\left[(a-\bar{a}(w))^2\right] \leq \frac{1}{\varepsilon_1}\frac{1-\delta}{\delta}\frac{\psi'(\bar{a}(w))^2}{\mathcal{I}(\bar{a}(w))}\frac{-\bar{F}''(w)}{2}. \tag{51}
$$

Since $\psi'(a)^2/\mathcal{I}(a)$ is Lipschitz continuous, there exists $\kappa > 0$ such that

$$
\begin{aligned}
& \frac{1-\delta}{\delta}\mathbb{E}^{\alpha}\left[\frac{\psi'(a)^2}{\mathcal{I}(a)}\right]\frac{\bar{F}''(w)}{2} + c(1-\delta)^{1+\varepsilon} \\
\leq \; & \frac{1-\delta}{\delta}\left(\frac{\psi'(\bar{a}(w))^2}{\mathcal{I}(\bar{a}(w))} + \left(\kappa\mathbb{E}^{\alpha}\left[|a-\bar{a}(w)|\right]\right)\right)\frac{\bar{F}''(w)}{2} + c(1-\delta)^{1+\varepsilon}.
\end{aligned}
$$

Note that

$$
\begin{aligned}
\frac{1-\delta}{\delta}\frac{-\bar{F}''(w)}{2}\kappa\mathbb{E}^{\alpha}\left[|a-\bar{a}(w)|\right] & \leq \frac{1-\delta}{\delta}\frac{-\bar{F}''(w)}{2}\kappa\sqrt{\mathbb{E}^{\alpha}\left[(a-\bar{a}(w))^2\right]} \\
& \leq \left(\frac{1-\delta}{\delta}\frac{-\bar{F}''(w)}{2}\right)^{\frac{3}{2}}\kappa\sqrt{\frac{\psi'(\bar{a}(w))^2}{\varepsilon_1\mathcal{I}(\bar{a}(w))}},
\end{aligned}
$$

where the first inequality follows by Cauchy-Schwarz, and the second follows by (51). Since this expression is of order $(1-\delta)^{3/2}$, there exists $\tilde{c} > 0$ such that, for sufficiently large δ, we have

$$\frac{1-\delta}{\delta}\mathbb{E}^{\alpha}\left[\frac{\psi'(a)^2}{\mathcal{I}(a)}\right]\frac{\bar{F}''(w)}{2}+c(1-\delta)^{1+\varepsilon}\leq\frac{1-\delta}{\delta}\frac{\psi'(\bar{a}(w))^2}{\mathcal{I}(\bar{a}(w))}\frac{\bar{F}''(w)}{2}+\tilde{c}(1-\delta)^{1+\varepsilon},$$

completing the proof.

G.3 Construction of $x_a(y)$

Lemma 25 *There exists $\bar{I}<\infty$ such that, for any sufficiently large $\delta<1$ and any $\bar{a}\in A$, there exists $x_{\bar{a}}:Y\to[-(1-\delta)^{-1/4},(1-\delta)^{-1/4}]$ satisfying*

$$\bar{a}\in\underset{a\in A}{\operatorname{argmax}}\int_y x_{\bar{a}}(y)\,p(y|a)\,dy-\psi(a),\tag{52}$$

$$\int_y x_{\bar{a}}(y)\,p(y|\bar{a})\,dy=0,\tag{53}$$

$$\int_y x_{\bar{a}}(y)^2\,p(y|\bar{a})\,dy\leq\frac{1}{\delta}\frac{\psi'(\bar{a})^2}{\mathcal{I}(\bar{a})},\qquad \textit{and}\tag{54}$$

$$\int_y \exp(\theta x(y|\bar{a}))\,p(y|\bar{a})\,dy\leq\exp\left(\theta^2\bar{I}\right).\tag{55}$$

Proof. Define, in turn,

$$\begin{aligned}\varphi_{\bar{a}} &= \frac{\mathcal{I}(\bar{a})}{\mathbb{E}\left[\mathbf{1}\left\{|\nu(y|\bar{a})|\leq(1-\delta)^{-\frac{1}{5}}\right\}\nu(y|\bar{a})^2\right]}, &(56)\\ \varepsilon_{\bar{a}} &= -\varphi_{\bar{a}}\frac{\psi'(\bar{a})}{\mathcal{I}(\bar{a})}\mathbb{E}\left[\mathbf{1}\left\{|\nu(y|\bar{a})|\leq(1-\delta)^{-\frac{1}{5}}\right\}\nu(y|\bar{a})\right],\qquad\text{and} &(57)\\ x_{\bar{a}}(y) &= \varphi_{\bar{a}}\frac{\psi'(\bar{a})}{\mathcal{I}(\bar{a})}\mathbf{1}\left\{|\nu(y|\bar{a})|\leq(1-\delta)^{-\frac{1}{5}}\right\}\nu(y|\bar{a})+\varepsilon_{\bar{a}}.\end{aligned}$$

Note that $\varphi_{\bar{a}}\geq 1$, since $\mathcal{I}(\bar{a})=\mathbb{E}\left[\nu(y|\bar{a})^2\right]$. We will prove that $x_{\bar{a}}(y)\in[-(1-\delta)^{-1/4},(1-\delta)^{-1/4}]$ for all y, and that (52)–(55) hold. We first establish that, for any sufficiently large $\delta<1$ and any $\bar{a}\in\left[0,\bar{A}\right]$, we have

$$|\varphi_{\bar{a}}-1|\leq\exp\left(-(1-\delta)^{\frac{1}{4}}\right)\qquad\text{and}\tag{58}$$

$$|\varepsilon_{\bar{a}}|\leq\frac{\psi'(\bar{a})}{\sqrt{\mathcal{I}(\bar{a})}}\exp\left(-(1-\delta)^{-\frac{1}{5}}\right).\tag{59}$$

Note that (58) and (59) immediately imply that $x_{\bar{a}}(y)\in[-(1-\delta)^{-1/4},(1-\delta)^{-1/4}]$.

For (58), note that

$$\begin{aligned} 0 \ \le\ \varphi_{\bar a} &= \frac{\mathcal{I}(\bar a)}{\mathcal{I}(\bar a) - \mathbb{E}\left[1\left\{|\nu(y|\bar a)| > (1-\delta)^{-\frac{1}{5}}\right\}\nu(y|\bar a)^2\right]} \\ &\le \frac{\mathcal{I}(\bar a)}{\mathcal{I}(\bar a) - \sqrt{\Pr\left(|\nu(y|\bar a)| > (1-\delta)^{-\frac{1}{5}}\right)\mathbb{E}\left[\nu(y|\bar a)^4\right]}}. \end{aligned}$$

where the second line follows by Cauchy-Schwarz. By Assumption 3(iii), we have

$$\int_y \exp(\theta\nu(y|\bar a))\,p(y|\bar a)\,dy \le \exp\left(\theta^2\mathcal{I}(\bar a)/2\right) \qquad \text{for all } \bar a \in A \text{ and } \theta \in \mathbb{R}. \tag{60}$$

Since $\mathcal{I}(\bar a)$ is uniformly bounded in $\bar a$ given Assumption 3(ii), there exists $c > 0$ such that, for all $\bar a \in A$, $\varphi_{\bar a}$ is bounded by

$$\frac{\mathcal{I}(\bar a)}{\mathcal{I}(\bar a) - \exp\left(-\frac{(1-\delta)^{-\frac{2}{5}}}{c}\right)16\Gamma(2)\,\mathcal{I}(\bar a)} = \frac{1}{1-\exp\left(-\frac{(1-\delta)^{-\frac{2}{5}}}{c}\right)16\Gamma(2)},$$

which implies (58).

For (59), note that

$$\begin{aligned} |\varepsilon_{\bar a}| &= \frac{\psi'(\bar a)}{\mathcal{I}(\bar a)}\left|\mathbb{E}\left[\varphi_{\bar a}\mathbf{1}\left\{|\nu(y|\bar a)| \le (1-\delta)^{-\frac{1}{5}}\right\}\nu(y|\bar a)\right]\right| \\ &\le \frac{\psi'(\bar a)}{\mathcal{I}(\bar a)}\left(\begin{array}{c} \left|\mathbb{E}\left[\mathbf{1}\left\{|\nu(y|\bar a)| > (1-\delta)^{-\frac{1}{5}}\right\}\nu(y|\bar a)\right]\right| \\ +(\varphi_{\bar a}-1)\left|\mathbb{E}\left[\mathbf{1}\left\{|\nu(y|\bar a)| \le (1-\delta)^{-\frac{1}{5}}\right\}\nu(y|\bar a)\right]\right| \end{array}\right) \end{aligned}$$

where the inequality is by $\mathbb{E}[\nu(y|\bar a)] = 0$ and the triangle inequality. As above, applying Cauchy-Schwarz and Assumption 3(ii)&(iii), we have

$$\frac{\psi'(\bar a)}{\mathcal{I}(\bar a)}\left|\mathbb{E}\left[\mathbf{1}\left\{|\nu(y|\bar a)| > (1-\delta)^{-\frac{1}{5}}\right\}\nu(y|\bar a)\right]\right| \le \exp\left(-\frac{(1-\delta)^{-\frac{2}{5}}}{c}\right)\frac{\psi'(\bar a)}{\sqrt{\mathcal{I}(\bar a)}}.$$

Again applying Cauchy-Schwarz, together with (58), we have

$$\frac{\psi'(\bar a)}{\mathcal{I}(\bar a)}(\varphi_{\bar a}-1)\left|\mathbb{E}\left[\mathbf{1}\left\{|\nu(y|\bar a)| \le (1-\delta)^{-\frac{1}{5}}\right\}\nu(y|\bar a)\right]\right| \le \frac{\psi'(\bar a)}{\sqrt{\mathcal{I}(\bar a)}}\exp\left(-(1-\delta)^{\frac{1}{4}}\right).$$

Taking $\delta < 1$ sufficiently large so that $\exp\left(-(1-\delta)^{-2/5}/c\right)+\exp\left(-(1-\delta)^{1/4}\right) \le \exp\left(-(1-\delta)^{1/5}\right)$, we have (59).

We now establish (52)–(55). Note that (53) follows directly from (57). For (52), for any $a \neq \bar{a}$, we have

$$
\begin{aligned}
& \mathbb{E}^{\bar{a}}\left[x_{\bar{a}}(y)\right]-\psi(\bar{a})-\left(\mathbb{E}^{a}\left[x_{\bar{a}}(y)\right]-\psi(a)\right) \\
= \quad & \psi(a)-\psi(\bar{a})-\varphi_{\bar{a}} \frac{\psi'(\bar{a})}{\mathcal{I}(\bar{a})} \int_{y} \mathbf{1}\left\{|\nu(y|\bar{a})| \leq(1-\delta)^{-\frac{1}{5}}\right\} \nu(y|\bar{a})\left(\begin{array}{c} p(y|a) \\ -p(y|\bar{a}) \end{array}\right) dy. \quad (61)
\end{aligned}
$$

We bound the second line as follows. For any $\gamma \in (0, \inf_{a\in\bar{A}} \psi''(a))$ and any $\bar{a} \in A$, we have

$$
\psi(a)-\psi(\bar{a}) \geq \psi'(a)(a-\bar{a})+\frac{\gamma}{2}(a-\bar{a})^{2};
$$

Taking a second-order Taylor expansion of $p(y|a)$ around $a=\bar{a}$, there exists $\hat{a} \in A$ such that

$$
\begin{aligned}
& \varphi_{\bar{a}} \frac{\psi'(\bar{a})}{\mathcal{I}(\bar{a})} \int_{y} \mathbf{1}\left\{|\nu(y|\bar{a})| \leq(1-\delta)^{-\frac{1}{5}}\right\} \nu(y|\bar{a})(p(y|a)-p(y|\bar{a})) dy \\
= \quad & (a-\bar{a}) \varphi_{\bar{a}} \frac{\psi'(\bar{a})}{\mathcal{I}(\bar{a})} \int_{y} \mathbf{1}\left\{|\nu(y|\bar{a})| \leq(1-\delta)^{-\frac{1}{5}}\right\} \nu(y|\bar{a})^{2} p(y|a) dy \\
& +\frac{(a-\bar{a})^{2}}{2} \varphi_{\bar{a}} \frac{\psi'(\bar{a})}{\mathcal{I}(\bar{a})} \int_{y} \mathbf{1}\left\{|\nu(y|\bar{a})| \leq(1-\delta)^{-\frac{1}{5}}\right\} \nu(y|\bar{a}) p_{aa}(y|\hat{a}) dy.
\end{aligned}
$$

Substituting (56), (61) is no less than $(a-\bar{a})^2/2$ multiplied by

$$
\gamma-\varphi_{\bar{a}} \frac{\psi'(\bar{a})}{\mathcal{I}(\bar{a})} \int_{y} \mathbf{1}\left\{|\nu(y|\bar{a})| \leq(1-\delta)^{-\frac{1}{5}}\right\} \nu(y|\bar{a}) p_{aa}(y|\hat{a}) dy.
$$

It remains to show that, for any sufficiently large $\delta < 1$ and any $\bar{a}, \hat{a}$, this expression is non-negative. Since $\psi'(\bar{a})/\mathcal{I}(\bar{a})$ is bounded, by (11) it suffices to show that

$$
\lim_{\delta \rightarrow 1} \sup_{\bar{a}, \hat{a}} \int_{y}\left(\varphi_{\bar{a}} \mathbf{1}\left\{|\nu(y|\bar{a})| \leq(1-\delta)^{-\frac{1}{5}}\right\}-1\right) \nu(y|\bar{a}) p_{aa}(y|\hat{a}) dy=0.
$$

In turn, it suffices to show that both

$$
\begin{aligned}
& \int_{y} \mathbf{1}\left\{|\nu(y|\bar{a})|>(1-\delta)^{-\frac{1}{5}}\right\} \nu(y|\bar{a}) p_{aa}(y|\hat{a}) dy \qquad \text{and} \\
& \left(\varphi_{\bar{a}}-1\right) \int_{y} \mathbf{1}\left\{|\nu(y|\bar{a})| \leq(1-\delta)^{-\frac{1}{5}}\right\} \nu(y|\bar{a}) p_{aa}(y|\hat{a}) dy
\end{aligned}
$$

converge to 0 as $\delta \rightarrow 1$, uniformly in $(\bar{a}, \hat{a})$.

By Cauchy-Schwarz, the first line is bounded by

$$\Pr\left(|\nu(y|\bar{a})| > (1-\delta)^{-\frac{1}{5}}\right)^{\frac{1}{4}} \left(\int_y \nu(y|\bar{a})^4 p(y|\bar{a})\,dy\right)^{\frac{1}{4}} \left(\int_y \left(\frac{p_{aa}(y|\hat{a})}{p(y|\bar{a})}\right)^2 p(y|\bar{a})\,dy\right)^{\frac{1}{2}}.$$

By (60), the first term of this product converges to 0 uniformly in $\bar{a}$ as $\delta \to 1$, and the second term is bounded uniformly in $\bar{a}$ given Assumptions 2(i) and (ii). Moreover, (12) ensures the last term is bounded uniformly in $(\bar{a}, \hat{a})$. So the entire product converges to 0 uniformly in $(\bar{a}, \hat{a})$.

Similarly, again by Cauchy-Schwarz, the second line is bounded by

$$(\varphi_{\bar{a}} - 1)\left(\int_y \nu(y|\bar{a})^2 p(y|\bar{a})\,dy\right)^{\frac{1}{2}} \left(\int_y \left(\frac{p_{aa}(y|\hat{a})}{p(y|\bar{a})}\right)^2 p(y|\bar{a})\,dy\right)^{\frac{1}{2}}.$$

By (58), the first term of this product converges to 0 uniformly in $\bar{a}$ as $\delta \to 1$; and, as above, the other terms are bounded uniformly in $(\bar{a}, \hat{a})$. The product thus converges to 0 uniformly in $(\bar{a}, \hat{a})$. This establishes (52).

We next establish (54). By construction, we have

$$\begin{aligned}\mathbb{E}\left[x_{\bar{a}}(y)^2\right] - \frac{1}{\delta}\frac{\psi'(\bar{a})^2}{\mathcal{I}(\bar{a})} &= \varphi_{\bar{a}}^2 \left(\frac{\psi'(\bar{a})}{\mathcal{I}(\bar{a})}\right)^2 \mathbb{E}\left[\mathbf{1}\left\{|\nu(y|\bar{a})| \leq (1-\delta)^{-\frac{1}{5}}\right\}\nu(y|\bar{a})^2\right] \\ &\quad + 2\varepsilon_{\bar{a}}\varphi_{\bar{a}}\frac{\psi'(\bar{a})}{\mathcal{I}(\bar{a})}\mathbb{E}\left[\mathbf{1}\left\{|\nu(y|\bar{a})| \leq (1-\delta)^{-\frac{1}{5}}\right\}\nu(y|\bar{a})\right] + \varepsilon_{\bar{a}}^2 - \frac{1}{\delta}\frac{\psi'(\bar{a})^2}{\mathcal{I}(\bar{a})}.\end{aligned}$$

By (56),

$$\left(\frac{\psi'(\bar{a})}{\mathcal{I}(\bar{a})}\right)^2 \varphi_{\bar{a}}\mathbb{E}\left[\mathbf{1}\left\{|\nu(y|\bar{a})| \leq (1-\delta)^{-\frac{1}{5}}\right\}\nu(y|\bar{a})^2\right] = \frac{\psi'(\bar{a})^2}{\mathcal{I}(\bar{a})}.$$

Thus, the above expression equals

$$\begin{aligned}&\varphi_{\bar{a}}(\varphi_{\bar{a}} - 1)\left(\frac{\psi'(\bar{a})}{\mathcal{I}(\bar{a})}\right)^2 \mathbb{E}\left[1\left\{|\nu(y|\bar{a})| \leq (1-\delta)^{-\frac{1}{5}}\right\}\nu(y|\bar{a})^2\right] \\ &+ 2\varepsilon_{\bar{a}}\varphi_{\bar{a}}\frac{\psi'(\bar{a})}{\mathcal{I}(\bar{a})}\mathbb{E}\left[1\left\{|\nu(y|\bar{a})| \leq (1-\delta)^{-\frac{1}{5}}\right\}\nu(y|\bar{a})\right] + \varepsilon_{\bar{a}}^2 - \frac{1-\delta}{\delta}\frac{\psi'(\bar{a})^2}{\mathcal{I}(\bar{a})}.\end{aligned}$$

Together with

$$\begin{aligned}
\left(\frac{\psi'(\bar{a})}{\mathcal{I}(\bar{a})}\right)^2 \mathbb{E}\left[1\left\{|\nu(y|\bar{a})| \leq (1-\delta)^{-\frac{1}{5}}\right\}\nu(y|\bar{a})^2\right] &\leq \frac{\psi'(\bar{a})^2}{\mathcal{I}(\bar{a})}, \\
\frac{\psi'(\bar{a})}{\mathcal{I}(\bar{a})}\mathbb{E}\left[1\left\{|\nu(y|\bar{a})| \leq (1-\delta)^{-\frac{1}{5}}\right\}\nu(y|\bar{a})\right] &\leq \frac{\psi'(\bar{a})}{\mathcal{I}(\bar{a})}\sqrt{\mathbb{E}\left[\nu(y|\bar{a})^2\right]} = \frac{\psi'(\bar{a})}{\sqrt{\mathcal{I}(\bar{a})}},
\end{aligned}$$

(58), and (59), this in turn is bounded by $\psi'(\bar{a})^2/\mathcal{I}(\bar{a})$ multiplied by

$$\varphi_{\bar{a}}\exp\left(-(1-\delta)^{\frac{1}{4}}\right) + 2\varphi_{\bar{a}}\exp\left(-(1-\delta)^{-\frac{1}{5}}\right) + \exp\left(-2(1-\delta)^{-\frac{1}{5}}\right) - \frac{1-\delta}{\delta},$$

which is non-positive uniformly in $\bar{a}$ for sufficiently large $\delta < 1$.

We finally establish (55). It suffices to show that, for any $I \geq \sup_{a\in A}\psi'(\bar{a})^2/\mathcal{I}(\bar{a})$, we have

$$\Pr(|x_{\bar{a}}(y)| \geq \lambda) \leq 2\exp\left(-\lambda^2/(4I)\right).$$

This is immediate if $\lambda \leq \sqrt{(\log 2)(4I)}$. Next, for any sufficiently large δ, $\bar{a} \in A$, and $\lambda \geq \sqrt{(\log 2)(4I)}$, we have $\lambda - \varepsilon_{\bar{a}} > 0$ and $-\lambda - \varepsilon_{\bar{a}} < 0$ by (59). Hence, we have

$$\begin{aligned}
\Pr(|x_{\bar{a}}(y)| \geq \lambda) &= \Pr\left(\frac{\psi'(\bar{a})}{\mathcal{I}(\bar{a})}\varphi_{\bar{a}}\mathbf{1}\left\{|\nu(y|\bar{a})| \leq (1-\delta)^{-\frac{1}{5}}\right\}\nu(y|\bar{a}) \leq -\lambda - \varepsilon_{\bar{a}}\right) \\
&\quad + \Pr\left(\frac{\psi'(\bar{a})}{\mathcal{I}(\bar{a})}\varphi_{\bar{a}}\mathbf{1}\left\{|\nu(y|\bar{a})| \leq (1-\delta)^{-\frac{1}{5}}\right\}\nu(y|\bar{a}) \geq \lambda - \varepsilon_{\bar{a}}\right) \\
&\leq \Pr\left(\frac{\psi'(\bar{a})}{\mathcal{I}(\bar{a})}\nu(y|\bar{a}) \leq \frac{-\lambda + \varepsilon_{\bar{a}}}{\varphi_{\bar{a}}}\right) + \Pr\left(\frac{\psi'(\bar{a})}{\mathcal{I}(\bar{a})}\nu(y|\bar{a}) \geq \frac{\lambda - \varepsilon_{\bar{a}}}{\varphi_{\bar{a}}}\right).
\end{aligned}$$

Since $\nu(y|\bar{a})$ is sub-Gaussian with variance-proxy $\mathcal{I}(\bar{a})$ by Assumption 3(iii), we have

$$\Pr\left(\frac{\psi'(\bar{a})}{\mathcal{I}(\bar{a})}\nu(y|\bar{a}) \leq \frac{-\lambda + \varepsilon_{\bar{a}}}{\varphi_{\bar{a}}}\right) + \Pr\left(\frac{\psi'(\bar{a})}{\mathcal{I}(\bar{a})}\nu(y|\bar{a}) \geq \frac{\lambda - \varepsilon_{\bar{a}}}{\varphi_{\bar{a}}}\right) \leq 2\exp\left(-\frac{\left(\frac{\lambda-\varepsilon_{\bar{a}}}{\varphi_{\bar{a}}}\right)^2}{2\frac{\psi'(\bar{a})^2}{\mathcal{I}(\bar{a})}}\right).$$

Finally, note that

$$\left(\frac{\left(\frac{\lambda-\varepsilon_{\bar{a}}}{\varphi_{\bar{a}}}\right)^2}{2\frac{\psi'(\bar{a})^2}{\mathcal{I}(\bar{a})}}\right) \Big/ \left(\frac{\lambda^2}{4I}\right) \geq 2\left(1 - \frac{2\varepsilon_{\bar{a}}}{\lambda}\right)\frac{1}{\varphi_{\bar{a}}^2}\frac{1}{\frac{\psi'(\bar{a})^2}{\mathcal{I}(\bar{a})}/I} \geq 2\left(1 - \frac{2\varepsilon_{\bar{a}}}{\sqrt{(\log 2)(4I)}}\right)\frac{1}{\varphi_{\bar{a}}^2},$$

which is greater than one uniformly in $\bar{a}$ for sufficiently large δ, by (58) and (59). We thus have

$$\Pr\left(\frac{\psi'(\bar{a})}{\mathcal{I}(\bar{a})}\nu(y|\bar{a}) \leq \frac{-\lambda - \varepsilon_{\bar{a}}}{\varphi_{\bar{a}}}\right) + \Pr\left(\frac{\psi'(\bar{a})}{\mathcal{I}(\bar{a})}\nu(y|\bar{a}) \geq \frac{\lambda - \varepsilon_{\bar{a}}}{\varphi_{\bar{a}}}\right) \leq 2\exp\left(-\frac{\lambda^2}{4I}\right),$$

as desired. ■

G.4 Equilibrium Verification

We verify that the contract defined in the main appendix, with $x_a(y)$ defined as in Lemma 25, satisfies incentive compatibility and promise keeping, and hence is an equilibrium.

Lemma 26 *For each h^t, we have*

$$\begin{aligned} a_t\left(h^t\right) &\in \operatorname*{argmax}_a (1-\delta)\left(u\left(c_t\left(h^t\right)\right) - \psi(a)\right) + \delta\mathbb{E}\left[w_{t+1}\left(h^{t+1}\right)|h^t, a\right], && (62) \\ w_t\left(h^t\right) &= (1-\delta)\left(u\left(c_t\left(h^t\right)\right) - \psi\left(a_t\left(h^t\right)\right)\right) + \delta\mathbb{E}\left[w_{t+1}\left(h^{t+1}\right)|h^t, a_t\left(h^t\right)\right]. && (63) \end{aligned}$$

Proof. The conclusion is immediate if h^t is irregular. If h^t is regular, then (62) follows from (52) since

$$\operatorname*{argmax}_a (1-\delta)\left(u\left(c_t\left(h^t\right)\right) - \psi(a)\right) + \delta\mathbb{E}\left[w_{t+1}\left(h^{t+1}\right)|h^t, a\right] = \operatorname*{argmax}_a \mathbb{E}\left[x_{\bar{a}(w_t(h^t))}(y_t)\,|h^t, a\right] - \psi(a).$$

Moreover, (63) holds as $u\left(\bar{c}\left(w_t\left(h^t\right)\right)\right) - \psi\left(\bar{a}\left(w_t\left(h^t\right)\right)\right) = w_t\left(h^t\right)$ by definition of $\bar{c}$ and $\bar{a}$, and $\mathbb{E}\left[w_{t+1}\left(h^{t+1}\right)|h^t, \bar{a}\left(w_t\left(h^t\right)\right)\right] = w_t\left(h^t\right)$ by (53). ■

G.5 Proof of Lemma 12

Let $\hat{H}^t = \left\{h^t : \left|w_t\left(h^t\right) - w\right| \leq (1-\delta)^{1/2-\varepsilon}\right\}$ be the set of regular period t histories. Let

$$x_{a_t}^{h^t}(y_t) = \begin{cases} x_{a_t}(y_t) & \text{if } h^t \in \hat{H}^t, \\ 0 & \text{if } h^t \notin \hat{H}^t, \end{cases} \qquad \text{and} \qquad X\left(h^t\right) = \frac{1-\delta}{\delta}\sum_{t'=1}^{t-1} x_{a_{t'}}^{h^{t'}}(y_{t'}),$$

where in the latter definition $h^{t'}$ is the period t' truncation of h^t. Note that $w_t\left(h^t\right) = w + X\left(h^t\right)$ for all t and h^t, and that $h^t \in \hat{H}^t$ iff $\left|X\left(h^t\right)\right| \leq (1-\delta)^{1/2-\varepsilon}$.

We first bound the weight on irregular histories under the equilibrium outcome μ.

Lemma 27 *For any sufficiently large $\delta < 1$, we have $(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\Pr^{\mu}\left(h^t \notin \hat{H}^t\right) \leq (1-\delta)^2$.*

Proof. Note that $x^{h^t}_{a_t}(y_t)$ is a sequence of martingale increments. Moreover, by (55),

$$\mathbb{E}^{\mu}\left[\exp\left(\theta\frac{1-\delta}{\delta}x^{h^t}_{a_t}(y_t)\right)|h^t,a_t\right] \leq \exp\left(\theta^2\left(\frac{1-\delta}{\delta}\right)^2\bar{I}\right) \qquad \text{for all } \theta, t, a_t.$$

Therefore, by Lemma 9,

$$\Pr^{\mu}\left(X\left(h^t\right)>x\right) \leq 2\exp\left(-\frac{x^2}{2\left(\frac{1-\delta}{\delta}\right)^2\bar{I}t}\right) \qquad \text{for all } x>0.$$

We can now apply Lemma 21 with $\vartheta=0$ to conclude that, for any sufficiently large δ,

$$(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\Pr^{\mu}\left(h^t\notin\hat{H}^t\right)=(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\Pr^{\mu}\left(\left|X\left(h^t\right)\right|>(1-\delta)^{1/2-\varepsilon}\right)\leq(1-\delta)^2.$$

■

Recall that $\varepsilon<1/8$. By Taylor expansion, since $3(1/2-\varepsilon)>1+\varepsilon$ and $|w_t(h^t)-w|\leq(1-\delta)^{1/2-\varepsilon}$ for all $h^t\in\hat{H}^t$, for any sufficiently large $\delta<1$ and any $h^t\in\hat{H}^t$, we have

$$\begin{aligned}
a_t\left(h^t\right)-c_t\left(h^t\right) &= \bar{a}\left(w_t\left(h^t\right)\right)-\bar{c}\left(w_t\left(h^t\right)\right)\\
&= \bar{F}\left(w_t\left(h^t\right)\right)\\
&\geq \bar{F}(w)+\bar{F}'(w)\left(w_t\left(h^t\right)-w\right)+\frac{\bar{F}''(w)}{2}\left(w_t\left(h^t\right)-w\right)^2-(1-\delta)^{1+\varepsilon}\\
&= \bar{F}(w)+\bar{F}'(w)\left(w_t\left(h^t\right)-w\right)+\frac{\bar{F}''(w)}{2}X\left(h^t\right)^2-(1-\delta)^{1+\varepsilon}.
\end{aligned}$$

At the same time, since $w_t(h^t)\in[0,2w]$, $a_t(h^t)-c_t(h^t)$ and $\bar{F}(w)+\bar{F}'(w)(w_t(h^t)-w)$ are bounded, and $\bar{F}''(w)\leq 0$, there exists $c_1>0$ such that, for any δ and $h^t\notin\hat{H}^t$, we have

$$a_t\left(h^t\right)-c_t\left(h^t\right)\geq\bar{F}(w)+\bar{F}'(w)\left(w_t\left(h^t\right)-w\right)+\frac{\bar{F}''(w)}{2}X\left(h^t\right)^2-c_1.$$

Combining these bounds, we have

$$\begin{aligned}
&(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\mathbb{E}^{\mu}\left[a_t\left(h^t\right)-c_t\left(h^t\right)\right]\\
\geq\;& \bar{F}(w)+\frac{\bar{F}''(w)}{2}(1-\delta)\sum_{t=2}^{\infty}\delta^{t-1}\int_h X\left(h^t\right)^2d\mu(h)\\
&-(1-\delta)^{1+\varepsilon}-c_1(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\Pr^{\mu}\left(h_t\notin\hat{H}^t\right).
\end{aligned} \tag{64}$$

Moreover, since $(x_{a_t}(y_t))_t$ is a sequence of martingale increments with variance bounded by (54), we have

$$\begin{aligned}
&(1-\delta)\sum_{t=2}^{\infty}\delta^{t-1}\sum_{h}\mu(h)\left(\frac{1-\delta}{\delta}\sum_{t'=1}^{t-1}x_{a_{t'}}(y_{t'})\right)^2 \\
= \ & \frac{1-\delta}{\delta}\left((1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\mathbb{E}^{a\sim\alpha_t}\left[\mathbb{E}^{y\sim p(y|a)}\left[x_a(y)^2\right]\right]\right) \leq \frac{1-\delta}{\delta^2}\mathbb{E}^{\alpha}\left[\frac{\psi'(a)^2}{\mathcal{I}(a)}\right],
\end{aligned}$$

Together with Lemma 27, (64) now implies that the principal's payoff is no less than

$$\bar{F}(w)+\frac{\bar{F}''(w)}{2}\frac{1-\delta}{\delta^2}\mathbb{E}^{\alpha}\left[\frac{\psi'(a)^2}{\mathcal{I}(a)}\right]-(1-\delta)^{1+\varepsilon}-(1-\delta)^2.$$

It remains to bound $\mathbb{E}^{\alpha}\left[\psi'(a)^2/\mathcal{I}(a)\right]$. Since $\psi'(a)^2/\mathcal{I}(a)$ is Lipschitz continuous, there exists $\kappa>0$ such that

$$\mathbb{E}^{\alpha}\left[\frac{\psi'(a)^2}{\mathcal{I}(a)}\right]-\frac{\psi'(\bar{a}(w))^2}{\mathcal{I}(\bar{a}(w))}\geq-\kappa\mathbb{E}^{\alpha}\left[|a-\bar{a}(w)|\right].$$

Since $\bar{a}(w)$ is continuously differentiable and $w\in[0,2w]$, there exists c_2 such that, for any t and $h^t\in\hat{H}^t$,

$$\left|a_t(h^t)-\bar{a}(w)\right|=\left|\bar{a}\left(w_t(h^t)\right)-\bar{a}(w)\right|\leq c_2\left|w_t(h^t)-w\right|\leq c_2(1-\delta)^{1/2-\varepsilon}.$$

Since $|a_t(h^t)-\bar{a}(w)|\leq\bar{A}$ for all t and h^t, we have

$$\begin{aligned}
\mathbb{E}^{\alpha}\left[|a-\bar{a}(w)|\right] &\leq (1-\delta)\sum_{t}\delta^{t-1}\left(\int_{h^t\in\hat{H}^t}c_2\left|w_t(h^t)-w\right|d\mu(h^t)+\Pr^{\mu}\left(h_t\notin\hat{H}^t\right)\bar{A}\right) \\
&\leq c_2(1-\delta)^{1/2-\varepsilon}+2\bar{A}(1-\delta)\sum_{t}\delta^{t-1}\Pr^{\mu}\left(h_t\notin\hat{H}^t\right) \\
&\leq c_2(1-\delta)^{1/2-\varepsilon}+2\bar{A}(1-\delta)^2,
\end{aligned}$$

where the third inequality follows from Lemma 27. We thus have

$$\mathbb{E}^{\alpha}\left[\frac{\psi'(a)^2}{\mathcal{I}(a)}\right]-\frac{\psi'(\bar{a}(w))^2}{\mathcal{I}(\bar{a}(w))}\geq-\kappa c_2(1-\delta)^{\frac{1}{2}-\varepsilon}-2\kappa\bar{A}(1-\delta)^2.$$

Therefore, by (64) (and using $(1-\delta)/\delta^2=(1-\delta)/\delta+O\left((1-\delta)^2\right)$), there exists c_3 such that, for any sufficiently large $\delta < 1$, the principal's payoff is no less than

$$\bar{F}(w) + \frac{1-\delta}{\delta} \frac{\psi'(\bar{a}(w))^2}{\mathcal{I}(\bar{a}(w))} \frac{\bar{F}''(w)}{2} - c_3 (1-\delta)^{1+\varepsilon},$$

completing the proof.