

Persuasion for the Long Run

James Best and Daniel Quigley*

July 18, 2022

Abstract

We examine persuasion when the sole source of credibility today is a desire to maintain a public record for accuracy. A long-run sender plays a cheap talk game with a sequence of short-run receivers, who observe some record of feedback about past accuracy. A geometric approach shows that when all feedback is public (as standard in repeated games), persuasion frequently requires inefficient on-path punishment—even if accuracy is monitored perfectly. If instead the record publishes coarse summary statistics (as is common online), any communication equilibrium the sender prefers to one-shot cheap talk—including Bayesian persuasion—can be supported without cost. (JEL C72, C73, D02, D82, D83)

Credible communication often relies on the desire to be believed not only today, but in the future as well. In many settings, like large anonymous markets, these long-run incentives depend in turn on the availability of public records. By publishing feedback about the quality of past communication, records make it possible to punish unsatisfactory advice with future incredulity. Yet such an endogenous source of commitment comes at a cost, as surplus-burning *punishments* must now occur whenever an inaccuracy is published.¹ This creates a trade-off new to persuasion: today’s approach to

*Best: Carnegie Mellon University Tepper School of Business (e-mail: jabest@andrew.cmu.edu). Quigley: Department of Economics and Exeter College, University of Oxford (e-mail: daniel.quigley@economics.ox.ac.uk). We thank Vince Crawford, Laura Doval, Johannes Hörner, Ian Jewitt, Emir Kamenica, Elliot Lipnowski, Laurent Mathevet, Margaret Meyer, John Hardman Moore, Alessandro Pavan, Doron Ravid, Maryam Saeedi, Ali Shourideh, Steven Tadelis, Ina Taneva, Adrien Vigier, Peyton Young and Bill Zame for their feedback.

¹Fudenberg and Levine (1994) show a folk theorem does not apply outside long-run relationships.

persuasion must be balanced against its likely effect on future credibility. Clearly, the process for publishing feedback must determine this trade-off. Hence, to analyze communication in these settings, we must understand the interaction between the individual persuasion problem and the design of the public record system itself.

With the rapid growth of online markets, managing these trade-offs is becoming ever more important. The approach of platforms such as eBay, Airbnb, and Upwork is to provide coarse public records. In particular, each promises its vendors a ‘badge’ if customer feedback meets the platform’s standards. Yet, customers are not allowed to view *all* relevant feedback.² We might worry that such systems weaken long-run incentives by hiding instances of feedback that could be used to punish participants. Perhaps surprisingly, the opposite is true: we show when platforms selectively pool feedback, myopic customers can be persuaded to incentivize suppliers more efficiently. Indeed, treating the choice of record as an information design problem, we show simple badge systems can support credible communication with almost no on-path punishment.

To explore the performance of public records as a source of commitment, we develop a model of long-run persuasion. A patient long-run sender (‘he’) plays a cheap talk game with a sequence of short-run receivers (each ‘she’). Each period a payoff-relevant state is independently drawn. The sender observes the state and sends a message to the receiver, who chooses an action. After she acts, a signal of the state (e.g., feedback) is realized. We say monitoring is perfect when the signal is noiseless, and imperfect otherwise. Repeated play of the game generates a history of messages, actions, and signals, which the sender always observes. Receivers may observe less: we say the public record is ‘complete’ if they observe the full history (standard in repeated games), and ‘incomplete’ otherwise. Our focus on short-run receivers means the sender can influence actions only via beliefs, allowing comparison with communication benchmarks like full information and Bayesian persuasion (Kamenica and Gentzkow (2011); KG, hereafter).

In section 2, we begin with an analysis of complete public records. Taking a belief-based approach, we apply the logic of Fudenberg et al. (1990) to recast the sender’s problem in terms of costly (one-shot) Bayesian persuasion. Theorem 1 shows that even in the best case of perfect monitoring, long-run incentives often fall short: a

²For example, eBay’s Top Rated badge requires sellers avoid “not as described” cases; these individual complaints are not published. Even where some individual reviews are available, they are rarely used: less than 1% of visitors click to see the details behind a rating (Nosko and Tadelis, 2015).

patient sender can achieve the Bayesian persuasion benchmark if and only if his optimal information structure under commitment is *partitional* (a deterministic mapping from states to messages, e.g., full information). Furthermore, Theorem 1 characterizes the limitations of long-run persuasion via the geometry of a *partitional value function*, \tilde{v} , which specifies at each prior the sender’s payoff from his optimal partitional information structure. We show long-run incentives are a *robust* substitute for commitment (i.e., the sender’s maximal payoff equals the KG benchmark at every prior) if and only if \tilde{v} is concave. As a result, the theorem provides a method for identifying the limitations of long-run persuasion without having to solve the Bayesian persuasion problem directly.

For some intuition, note that deviations from partitional information structures are easily identified ex post when monitoring is perfect. Hence, they can be enforced via off-path punishments. By contrast, non-partitional information involves mixing between different messages in at least one state. Such mixing requires indifference, which in turn requires (surplus-burning) on-path punishments.³ So if long-run incentives are a robust substitute for commitment, the KG payoff must be attained with some partitional information structure for every prior and \tilde{v} must inherit the concavity of the Bayesian persuasion benchmark. Perhaps less obvious, the converse is also true.

We examine some applied consequences of this geometry for long-run persuasion. Proposition 3 establishes some bad news when the set of actions is finite: if the sender can benefit from persuasion, then long-run persuasion almost always fails to attain the KG payoff. Proposition 4 evaluates the efficacy of certification using *fixed* (prior-independent) partitions, such as grading students or rating debt. We provide necessary and sufficient conditions under which such communication always attains KG payoffs.

Fundamentally, non-partitional communication is costly because credibility requires a threat of punishment, and mixed strategies cause this punishment to occur on path. Moreover, when monitoring is imperfect, even partitional information structures involve on-path punishment, and the equilibrium payoff set contracts (Fudenberg and Levine, 1994). However, in section 3 we present our second main insight: by pooling histories we can persuade receivers to discipline the sender more efficiently, benefiting sender and receivers alike. Applying information design to the public record, we show credibility can be supported with almost no on-path punishment.

³Mathevet et al. (2019) makes the same initial observation. They then take a different approach from us, studying a reputational-types model in which the sender is committed with some probability.

To fix ideas, consider an online seller advertising products of differing qualities at a fixed price. To secure a sale, the seller must persuade a customer that the product’s expected quality is sufficiently high; if he does, the customer leaves a noisy review of her purchase (i.e., monitoring is imperfect). Consider for a moment an equilibrium in which at some time t two types of feedback history may arise: those at which the seller is incentivized to recommend only high quality products, and those at which he is not. If the public record is complete, both types of history will involve costs in equilibrium: the seller expects future punishments in the former, and cannot sell in the latter.

Suppose a designer could pool these feedback histories from the perspective of the t^{th} customer only. If the seller adopts the same (now private) strategy, this customer may now be persuaded to follow a ‘buy’ recommendation at *both* histories (given her information, she will account for the seller’s *expected* play across these histories). Moreover, because all feedback is released after t , future customers’ ability to discipline the the seller’s private strategy at t is not diminished. Hence, the t^{th} customer can be persuaded to buy more often, but with less on-path punishment. Clearly, the seller is better off. More surprisingly, we show the customers may be better off too—the reduction in punishments can more than compensate for the ex ante cost of occasionally buying from an uninformative seller.

In section 3.2 we explore the scope for the design of coarse records to improve payoffs. Theorem 2 shows that under a broad set of conditions, a designer (‘the platform’) can employ a *simple badge system* (SBS) to almost costlessly implement any information structure that the sender prefers to one-shot cheap talk. Hence, the scope for gains can be large: in our online trade example, it covers all feasible communication. In a SBS, the platform collects and periodically evaluates receivers’ feedback against a set of chosen standards. These might, for example, require a seller to keep the rate of complaints below a threshold. The platform awards the sender with a public badge if and only if its standards are met. While the sender observes all feedback, incoming receivers can observe *only* the badge; even the dates of evaluation are hidden from them.⁴

To implement a target information structure, the platform uses long evaluations and chooses standards strict enough that communication ‘far’ from the target loses the badge with probability close to one, but lenient enough that communication ‘near’ the

⁴Similar to Kremer et al. (2014) we model this as hiding information about uncertain arrival time within an evaluation phase. We discuss various foundations for this assumption in section 4.

target very likely retains it. This design admits an equilibrium which reduces on-path punishments by disciplining average behavior only, as in Radner (1985). The sender communicates near the target *on average* to avoid being punished by receivers for losing the badge. Consequently, when he has the badge, receivers’ beliefs are determined by this average alone. However, this necessarily abandons incentivizing some interactions: if the sender has already done enough to retain the badge, he has no incentive to communicate on target. Clearly, if a receiver knew she were at such a history she would not find him credible, and so her strategy would not be a best response. The SBS solves this problem by pooling histories to persuade receivers to adopt this more efficient strategy. Indeed, such persuasion is at the heart of any incomplete record system that improves equilibrium payoffs (relative to the complete record case).⁵

These results are particularly relevant to the flourishing of online trade.⁶ As Tadelis (2016) points out, this success has relied largely on the development of feedback and reputation systems, which have been carefully refined by platforms since their inception and now often take the form of badges.⁷ For instance, eBay periodically assesses its sellers and awards a Top Rated badge to those who have attracted disputes on no more than 0.5% of their sales. To avoid disputes, sellers are advised to “describe (each) item accurately”. Airbnb and Upwork also provide badges that require accurate communication about availability, among other things. These badges all depend on hiding historical data from other users, such as the the rate of disputes (eBay) or individual ratings (Airbnb). Finally, eBay and Upwork provide vendors with portals to *privately* track their feedback and manage any risk of losing their badge.

One might worry that in reality platforms could face issues committing to a simple badge system. In section 4 we discuss this issue in light of an alternative (sender-blind) review system. In doing so, we also formally identify the theoretical and applied differences between public record systems and approaches that ‘reuse punishments’, as in Abreu et al. (1991) and Fuchs (2007). We also extend our analysis to evaluate review systems that cannot hide evaluation dates or acquire feedback, and those that face other forms of moral hazard. Finally, we discuss related literature in section 5 and conclude in section 6. Proofs are contained in the appendices.

⁵That is, relative to the payoff bounds arising from the logic of Fudenberg and Levine (1994).

⁶Having grown from 0.5% to 14% of retail sales in the past two decades (US Bureau of the Census).

⁷Hui et al. (2016) find that eBay’s Top Rated badge both confers significant reputational advantages to sellers and motivates those in danger of losing the badge to improve their behavior.

1 Model

We consider a general repeated cheap talk game between a long-run sender ('he') with discount rate $\delta < 1$ and a sequence of short-run receivers (each 'she'). At each time $t = 1, 2, \dots$, the current state of the world θ_t is drawn randomly from a common prior μ_0 over a finite set Θ ; θ is realized with probability $\mu_{0,\theta}$. The sender privately observes θ_t and sends a message $m_t \in M$ to the receiver. We allow M to be an uncountable Polish space endowed with the Borel σ -algebra. However, little is lost if the reader prefers to treat M as large but finite (see Proposition 1). On observing m_t , the receiver chooses an action a_t from a compact, metrizable set A . The sender's and receivers' respective stage payoffs are $v(a)$ and $u(\theta, a)$, each continuous in a . We assume the sender cares only about a so his short-term incentive is to say whatever induces his favored action.⁸ At the end of the stage game, a noisy signal ω_t may be drawn from a finite set Ω . For all but at most one action a' , we assume ω is drawn from a conditional distribution $p(\omega_t | \theta_t)$, where each (θ_t, ω_t) is history-independent. To accommodate natural applications (see example 1), we allow a' (interpreted as 'inaction') to induce no signal. Except for inaction, monitoring therefore takes a product structure (Fudenberg et al. (1994)). We say monitoring is **perfect** if ω_t identifies θ_t for any $a_t \in A$, and **imperfect** otherwise.

On entering period t , the sender observes past $\theta^t = (\theta_1, \dots, \theta_{t-1})$ and the **history of feedback**, $\underline{h}_t = (m_\tau, a_\tau, \omega_\tau)_{\tau=1}^{t-1}$, with $\underline{h}_1 = \emptyset$. Denote the sender's (private) history by $h_t = (\theta^t, \underline{h}_t)$, and let the corresponding sets of histories at t be $\underline{\mathcal{H}}_t$ and \mathcal{H}_t , respectively. We write $\underline{\mathcal{H}}, \mathcal{H}$ for the sets of all such histories. The receiver at t only observes a public record r_t , determined by a sequence of functions, $r_t : \underline{\mathcal{H}}_t \rightarrow \mathcal{R}_t$, $t = 1, 2, \dots$, where each \mathcal{R}_t is the set of possible public records at t and $\mathcal{R} = \bigcup_{t=1}^{\infty} \mathcal{R}_t$. Let \mathcal{R}^t be the set of time t sequences $r^t = (r_1, \dots, r_t)$. Our analysis compares different types of record: a **complete public record** sets $\mathcal{R}_t = \underline{\mathcal{H}}_t$ and $r_t = \underline{h}_t$ for all t .^{9,10} We call other (non-invertible) forms of public record **incomplete**; we discuss them in detail in section 3. Throughout, we allow players access to a public randomization device.

The sender's payoff in the repeated game is $\sum_{\tau=1}^{\infty} \delta^{\tau-1} v(a_\tau)$. To make welfare comparisons with static benchmarks, we occasionally discount receivers' payoffs at the same rate. Strategies for the sender $\sigma : \mathcal{H} \times \Theta \rightarrow \Delta M$, receiver $\rho : \mathcal{R} \times M \rightarrow \Delta A$, and belief

⁸This is largely expositional: Theorems 1 and 2 can be extended to state-dependent v .

⁹Since M and therefore each $\underline{\mathcal{H}}_t$ need not be finite, each \mathcal{R}_t may also be infinite in principle.

¹⁰When clear from context, we may use the same notation for a function and its output (as for r_t).

systems $\mu : \mathcal{R} \times M \rightarrow \Delta\Theta$ are measurable functions, defined on the respective histories.^{11,12} A profile $\langle \sigma, \rho, \mu \rangle$ constitutes a perfect Bayesian equilibrium (PBE) if σ is a best response for the sender at each (h_t, θ_t) , ρ a best response for the receiver at each (r_t, m_t) , and beliefs are derived from Bayes rule where possible (see Appendix A for formalities).¹³ A public PBE (PPBE) is a PBE in which the sender's strategy depends only on r_t and θ_t . We illustrate our main results with the following example:

Example 1 *A long-run seller serves a sequence of anonymous customers. Each period, he has a product of quality $\theta \in \{l < 0, h = 1\}$, where $\Pr[\theta = h] = \mu_0$, and makes a claim $m \in \{\text{'like new'}, \text{'used'}\}$. If a customer 'Refuses' the product, she and the seller get payoff 0. If she 'Buys' she gets θ , while he gets 1. Feedback $\omega \in \{b, g, \emptyset\}$ is distributed according to $\Pr(g | h, \text{Buy}) = \Pr(b | l, \text{Buy}) = p > \frac{1}{2}$ and $\Pr(\emptyset | h, \text{Refuse}) = 1$. Hence, g , b , and \emptyset can be interpreted as good, bad, and no review, respectively.*

Benchmarks for Long-Run Persuasion

In any stage game, the sender's strategy induces an **information structure**: a distribution $\lambda \in \Delta\Delta\Theta$ over the receiver's posterior beliefs. When the sender can commit to information structures, KG shows it is without loss of generality for average payoffs to (i) restrict strategies to a choice of information structure from the *Bayes plausible* set $\Lambda(\mu_0) = \{\lambda : \mathbb{E}_\lambda[\mu] = \mu_0\}$, with $|\text{supp } \lambda| \leq |\Theta| + 2$; and (ii) break receiver indifference in favor of the sender. When discussing commitment benchmarks, we also adopt these properties. Let $\lambda(\mu)$ and $\lambda(\mu | \theta)$ denote the marginal and conditional probabilities that the receiver's induced posterior is μ . With slight abuse of notation, we write the sender's payoff given posterior μ as $v(\mu)$, and $v(\lambda) := \mathbb{E}_\lambda[v(\mu)]$ for its mixed extension; $u(\mu)$, $u(\lambda)$ are used analogously for receivers; we denote their pure best response by $a(\mu)$. The set of payoff pairs attainable with commitment is

$$\mathcal{CS} := \left\{ (u, v) : u = u(\lambda), v = v(\lambda), \text{ for some } \lambda \in \Lambda(\mu_0) \right\}.$$

¹¹Since r_t is a function of \underline{h}_t , we can write σ as a reduced-form function of (h_t, θ_t) only.

¹² ΔX is the space of probability measures on X with the weak* topology, unless stated otherwise.

¹³Formally we require beliefs are given by an appropriate probability kernel. See Appendix A. As only induced beliefs about θ_t matter, we need not explicitly reference beliefs over \mathcal{H} .

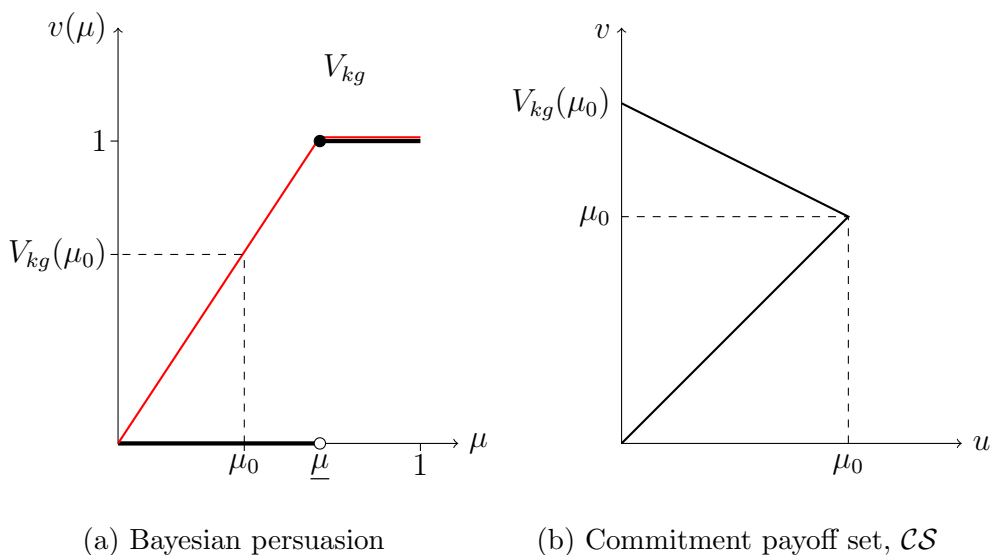
The **Bayesian persuasion payoff** is the sender’s maximal payoff among those in \mathcal{CS} :

$$V_{kg}(\mu_0) := \max_{\lambda \in \Lambda(\mu_0)} v(\lambda). \quad (\text{KG})$$

KG show $V_{kg} = \text{cav } v$, the concave envelope of the (interim) value function $v(\mu)$. We call any solution to (KG) **Bayesian persuasion**, denoted $\lambda^{kg} \in \Lambda^{kg}(\mu_0)$. Of course, the receiver’s optimum is attained by full information.

Figure 1 illustrates these concepts for example 1. The customer buys if and only if her posterior is $\mu \geq \underline{\mu} := \frac{|l|}{1+|l|}$. Hence, $v(\mu)$ is the thick black line in Fig. 1a, $V_{kg}(\mu)$ the red line, and $V_{kg}(\mu_0) = \frac{\mu_0}{\underline{\mu}}$. Notice that if $\mu_0 \geq \underline{\mu}$, then λ^{kg} will be uninformative. Otherwise, λ^{kg} splits the prior to posteriors 0 and $\underline{\mu}$. This is equivalent to the seller reporting ‘like new’ with conditional probabilities $\lambda(\underline{\mu}|\theta = h) = 1$ and $\lambda(\underline{\mu}|\theta = l) = \frac{\mu_0}{(1-\mu_0)|l|}$. The customer optimum is full information, with payoffs (μ_0, μ_0) in Fig. 1b.

Figure 1: Bayesian persuasion in example 1



Finally, we assume that informative Bayesian persuasion cannot be replicated in a one-shot cheap talk equilibrium of the stage game. Defining $V_{ct}(\mu)$ as the sender’s maximal equilibrium payoff in the one-shot version of the game, our assumption is¹⁴

Assumption 1 For every $\mu \in \Delta\Theta$, $V_{kg}(\mu) > v(\mu)$ implies $V_{kg}(\mu) > V_{ct}(\mu)$.

¹⁴Lipnowski and Ravid (2020) show V_{ct} is well-defined and characterize it geometrically for state-independent v : V_{ct} is the smallest quasi-concave function everywhere weakly greater than v .

Assumption 1 implies that whenever λ^{kg} is informative the sender is not indifferent between all posteriors in $\text{supp } \lambda^{kg}$, but imposes no such requirement at priors for which λ^{kg} is uninformative. For this reason, it is not very restrictive.¹⁵

2 Persuasion with Complete Public Records

We begin by investigating how well repetition substitutes for ex ante commitment in persuasion with complete public records—the standard setting in repeated games. For complete records, we denote the set of (discounted average) PBE payoff pairs by $\mathcal{E}(\mu_0, \delta)$.

2.1 Long-Run Persuasion as Costly Persuasion

Our setting features repeated play and no exogenous commitment. Nonetheless, we show an analogue of KG’s belief-based approach still applies. This helps us reduce a dynamic problem down to a static costly persuasion problem. Such an approach is convenient: it facilitates comparison with the Bayesian persuasion benchmark and allows us to geometrically characterize the limitations of long-run persuasion.

To see this, note first that the game described in section 1 has an associated ‘direct game’, identical except that the message space is now the set of recommended beliefs $\Delta\Theta$. In this direct game, we identify the sender’s (mixed) strategy with a function $\lambda^d : \underline{\mathcal{H}} \rightarrow \Delta\Delta\Theta$, specifying an information structure at each public history.¹⁶ Receivers’ strategies and beliefs, denoted ρ^d and μ^d respectively, are defined analogously to those in section 1 (see Appendix A). In a direct game, there is no reason for receivers’ equilibrium beliefs to coincide with m_t or for the sender’s chosen information structures to even be Bayes plausible. Still, we introduce a class of equilibria in which these features arise:

Definition 1 A *direct equilibrium* is a PPBE $\langle \lambda^d, \rho^d, \mu^d \rangle$ of the direct game, in which for all $\underline{h}_t \in \underline{\mathcal{H}}$: (i) $\lambda^d(\underline{h}_t) \in \Lambda(\mu_0)$, and (ii) $\mu^d(\underline{h}_t, m_t) = m_t, \forall m_t \in \text{supp } \lambda(\underline{h}_t)$.

Clearly, the direct equilibrium (DE) concept applies only to the direct game, and our analysis takes this for granted from here. DE is the analogue of the belief-based ap-

¹⁵Functions violating Assumption 1 are non-generic (Best and Quigley, 2017) .

¹⁶The marginals $\lambda^d(\underline{h}_t)$ are generated by a public strategy $\sigma : \underline{\mathcal{H}} \times \Theta \rightarrow \Delta\Delta\Theta$. See Appendix A.

proach of KG to long-run persuasion: condition (i) requires that the sender's equilibrium strategy induces a Bayes plausible distribution over recommended beliefs at each public history of the repeated game, and condition (ii) requires that the receivers' beliefs match the sender's recommendations on path. Unlike KG, DE additionally imposes equilibrium incentive constraints, which capture, *inter alia*, the sender's lack of commitment.¹⁷ Moreover, in DE the sender communicates only about θ_t , but not his private history. Nonetheless, the next two results show the belief-based approach is still without loss:¹⁸

Lemma 1 *A PBE of the game induces a collection of marginal distributions \mathbb{P}_t over $(\theta_t, a_t, \omega_t)$, $t = 1, 2, \dots$, if and only if there is a DE that induces the same marginals.*

Lemma 1 tells us that we can characterize the set of PBE outcomes by focusing on DE alone. As payoffs are additively separable over time, the relevant outcomes are the marginal distributions of the contemporaneous variables at each time t . The proof shows how to relabel on-path messages so that they can be interpreted as recommended beliefs without introducing measurability issues. The next proposition applies the recursive techniques of Abreu et al. (1990) to show that the DE payoff set can be fully characterized by information structures with bounded support.

Proposition 1 *$\mathcal{E}(\mu_0, \delta)$ is compact, convex, and increasing in δ in the set inclusion order. If $(u, v) \in \mathcal{E}(\mu_0, \delta)$, there is a DE in which $|\text{supp } \lambda(\underline{h}_t)| \leq |\Theta| + 1$, $\forall \underline{h}_t \in \underline{\mathcal{H}}$, and payoffs are (u, v) .*

As off-path messages are easily disciplined, no new issues arise from our sequential stage game and the set of DE payoffs coincides with the set of perfect public equilibrium payoffs, to which recursive tools apply. That $|\text{supp } \lambda(\underline{h}_t)|$ can be uniformly bounded is a consequence of the convexity of $\mathcal{E}(\mu_0, \delta)$ and the finiteness of Θ . As $\mathcal{E}(\mu_0, \delta)$ is compact, a sender-optimal DE exists, with payoff:

$$V_{de}(\mu_0, \delta) = \max\{v : (u, v) \in \mathcal{E}(\mu_0, \delta)\}. \quad (1)$$

Trivially, V_{de} is bounded above by V_{kg} . After all, the sender is constrained not only by Bayes plausibility but also by incentive compatibility at each history. To establish a

¹⁷As deviations from $\lambda^d(\underline{h}_t)$ need not live in $\Lambda(\mu_0)$, the sender's incentives are the same as in PBE.

¹⁸Lipnowski and Ravid (2020) have a similar result in the context of one-shot cheap talk.

more meaningful bound, define $\underline{v}_\theta(\lambda) := \min_{\mu \in \text{supp } \lambda | \theta} v(\mu)$ as the sender's worst payoff induced in state θ by a λ with finite support. We can now apply Fudenberg et al. (1990) to DE, reducing the dynamic problem to one of costly (one-shot) Bayesian persuasion.¹⁹

Proposition 2 (Fudenberg et al., 1990) *The sender's optimal continuation payoff in any long-run persuasion game is bounded above by*

$$\max_{\lambda \in \Lambda(\mu_0)} v(\lambda) - c(\lambda), \quad (\text{CP})$$

where $c(\lambda) := \mathbb{E}_\lambda [v(m) - \underline{v}_\theta(\lambda)] \geq 0$. If monitoring is perfect, then there exists $\underline{\delta} < 1$ such that $V_{de}(\mu_0, \delta)$ is equal to (CP) for all $\delta \in [\underline{\delta}, 1)$.

The sender's equilibrium payoffs can never exceed those he would get if he could commit to an information structure but was only able to receive his worst payoff in each state. In order to persuade any receiver, the sender's equilibrium information structure must be credible: he must prefer it to any deviation. Hence, if his optimal information structure involves mixing between messages in some state θ , he must be indifferent between them in that state. In the sender's best equilibrium, this indifference requires that each message m induces an on-path punishment $v(m) - \underline{v}_\theta(\lambda)$. This exactly offsets the stage gains of each message relative to the worst, implying an expected punishment cost of $c(\lambda)$ and pinning down the upper-bound. Figure 2a illustrates this for example 1 with $\mu_0 < \underline{\mu}$; the equilibrium payoff set is a strict subset of \mathcal{CS} . Here, the seller can do no better than an equilibrium in which he provides full information on path.

2.2 The Limits of Long-Run Persuasion

Expression (CP) provides the best case for repetition as a substitute for commitment. Hence, in this section, we assume monitoring is perfect and δ sufficiently large.²⁰ Under these conditions, (CP) suggests **partitional** information structures are particularly relevant. An information structure is said to be partitional if it is generated by a pure strategy. For any such pure strategy there is a corresponding partition of Θ , $\mathcal{P} = \{P_1, \dots, P_k\}$ for which each message represents a claim that ' θ is in P_i '. Given μ_0 , communication

¹⁹Since it is a direct application, we omit the proof. Mathevet et al. (2019) also identifies this bound.

²⁰As we consider the patient limit here, we drop dependence of V_{de} on δ in the notation.

over \mathcal{P} induces an information structure $\lambda^{\mathcal{P}}$ (the obvious dependence on μ_0 suppressed) supported on $\{\mu_i\}_{i=1}^k$, where $\mu_i = (\mu_{i,\theta})_{\theta \in \Theta}$ denotes the posterior distribution associated with learning $\{\theta \in P_i\}$.²¹ Hence, μ_i lies in $\Delta P_i := \{\mu \in \Delta\Theta : \mu_\theta = 0, \forall \theta \notin P_i\}$.²² As μ_i is induced whenever $\theta \in P_i$, it occurs with probability $\lambda_i := \lambda^{\mathcal{P}}(\mu_i) = \sum_{\theta \in P_i} \mu_{0,\theta}$, and the sender's associated stage payoff is $v^{\mathcal{P}}(\mu_0) := \sum_{i=1}^k \lambda_i v(\mu_i)$. Partitional information structures have the advantage of requiring no on-path punishment, i.e., $c(\lambda^{\mathcal{P}}) = 0$. Letting \mathcal{P}^Θ be the set of all partitions of Θ , we define the sender's maximal payoff from such communication:

Definition 2 The *partitional value function* $\tilde{v} : \Delta\Theta \rightarrow \mathbb{R}$ is defined by

$$\tilde{v}(\mu) = \max_{\mathcal{P} \in \mathcal{P}^\Theta} v^{\mathcal{P}}(\mu). \quad (2)$$

Figure 2: Long-run persuasion in example 1

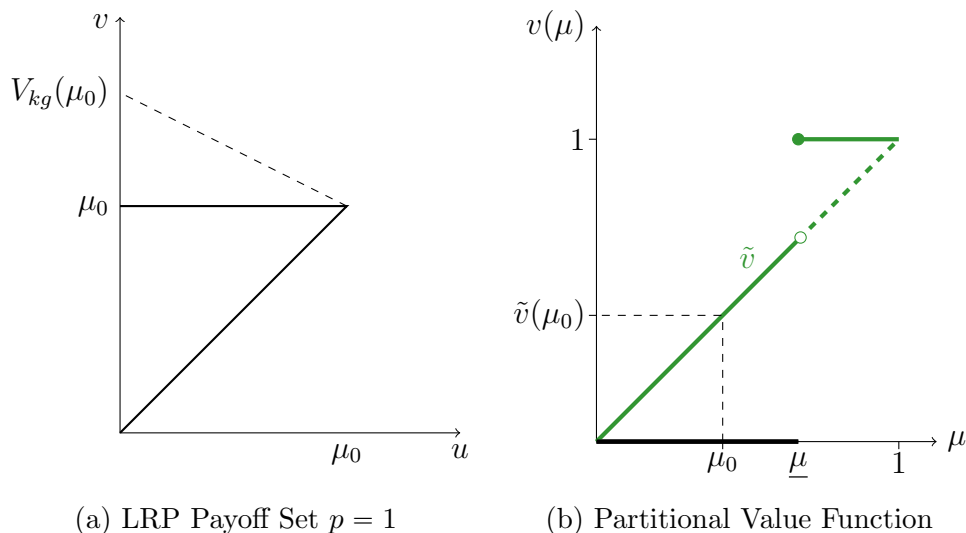


Figure 2b illustrates \tilde{v} in the context of example 1. Of course, $\tilde{v} \leq V_{de} \leq V_{kg}$. This simple fact helps to characterize the limitations of long-run persuasion.

Theorem 1 For a fixed prior, $V_{de}(\mu_0) = V_{kg}(\mu_0)$ if and only if $V_{kg}(\mu_0) = \tilde{v}(\mu_0)$. Moreover, $V_{de}(\mu_0) = V_{kg}(\mu_0)$ for all priors if and only if \tilde{v} is concave.

Theorem 1 first establishes that we can attain the Bayesian persuasion payoff at some prior μ_0 if and only if there exists some partitional $\lambda^{kg} \in \Lambda^{kg}(\mu_0)$. Hence, a

²¹Analogous to $\mu_{0,\theta}$, $\mu_{i,\theta}$ is the probability of state θ under the belief μ_i .

²²This definition abuses notation but is convenient as it implies $\Delta\Theta$ and ΔP_i live in the same space.

direct comparison of \tilde{v} and V_{kg} identifies the limitations of long-run persuasion. But beyond direct comparison, it also provides an alternative approach to evaluating these limitations based on the geometry of \tilde{v} alone. Specifically, the second condition implies that $V_{de} = V_{kg}$ if and only if \tilde{v} satisfies Jensen's inequality for all pairs in the simplex.²³ This can be beneficial: as Propositions 3 and 4 illustrate, evaluating the payoffs of binary lotteries can be easier than direct comparison of value functions, not least because finding V_{kg} can involve optimization over a much larger space than that of binary lotteries (Lipnowski and Mathevet (2017a)). Further, the theorem holds independent interest: in settings where knowledge of prior beliefs is scant, it identifies when long-run incentives robustly substitute for commitment.

For the first equivalence, sufficiency is immediate. To see why $\tilde{v}(\mu_0) = V_{kg}(\mu_0)$ is also necessary, consider the case in which $\lambda^{kg} \in \Lambda(\mu_0)$ uniquely maximizes the sender's payoff.²⁴ Now suppose that λ^{kg} is not partitional but $V_{de}(\mu_0) = V_{kg}(\mu_0)$. This implies that in some state θ , λ^{kg} induces at least two beliefs μ' , μ'' with identical payoffs (otherwise, a punishment would be associated with one of them). Since λ^{kg} is uniquely optimal, v must be strictly greater at μ' and μ'' than at any μ on the line segment connecting them: $v(\mu) < v(\mu') = v(\mu'')$. Yet, for any prior μ on this line, a splitting to $\{\mu', \mu''\}$ yields payoff $V_{kg}(\mu) > v(\mu)$ and is a one-shot cheap talk equilibrium. Hence $V_{ct}(\mu) = V_{kg}(\mu) > v(\mu)$, violating Assumption 1.

The second equivalence follows quickly from the first. Indeed, the necessity of concave \tilde{v} is virtually a restatement of the first condition. Sufficiency follows from two observations: first, since silence is itself trivially partitional, \tilde{v} must satisfy $v \leq \tilde{v} \leq V_{kg}$; second, V_{kg} is by definition the smallest concave function that is everywhere weakly greater than v . Hence, if \tilde{v} is concave, it must everywhere equal V_{kg} .

In our leading example, it is clear from Figure 2b that \tilde{v} is not concave. Moreover, $V_{de} < V_{kg}$ at any prior where the sender can benefit from persuasion, i.e., $\mu_0 < \underline{\mu}$. However, it is not hard to find conditions under which $V_{de} = V_{kg}$ for all priors. For instance, \tilde{v} is trivially concave when v is either concave or convex: in each case a single partition (the coarsest and finest, respectively) attains V_{kg} everywhere. We leverage Theorem 1 to ask how these two observations generalize. First, what are the limitations

²³In KG, the sender can benefit from persuasion if and only if $\text{cav } v > v$; relatedly, a corollary of Theorem 1 is that the sender would benefit from ex ante commitment power if and only if $\text{cav } \tilde{v} > \tilde{v}$.

²⁴The argument is similar in spirit for multiple maximizers, albeit a little more involved.

of long-run persuasion when actions are finite? Second, under what conditions does communication over a single partition always attain Bayesian persuasion payoffs?

We answer the first question under the additional assumption that receivers have a strict preference ordering over actions in A for each $\theta \in \Theta$: if $a \neq a'$ then $u(\theta, a) \neq u(\theta, a')$. We call these preferences **conditionally strict**.²⁵

Proposition 3 *Suppose A is finite and receivers have conditionally strict preferences. Then, μ_0 satisfies one of the following, (Lebesgue) a.e. on $\Delta\Theta$:*

1. $\mu_0 \in \arg \max_{\Delta\Theta} v(\mu)$; or
2. $V_{kg}(\mu_0) > \tilde{v}(\mu_0)$.

When $\mu_0 \in \arg \max v(\mu)$, long-run persuasion trivially attains the KG benchmark. Yet, as the receiver will take the sender's favorite action anyway, he cannot benefit from persuasion. By contrast, when the sender can actually benefit from persuasion, long-run persuasion almost always falls short of the KG benchmark. In short, long-run incentives are typically a costly substitute for commitment whenever actions are finite.²⁶

To prove this, we show that for almost all $\mu \notin \arg \max v$ there exists a splitting of μ to $\{\mu', \mu''\}$ for which the partitional value function \tilde{v} violates Jensen's inequality. Let the optimal partition at μ be \mathcal{P} . Specifically, we show μ'' can be chosen so that $\tilde{v}(\mu'')$ exceeds $v^{\mathcal{P}}(\mu)$, and μ' picked close enough to μ such that the payoff from disclosing $\{\theta \in P_i\}$ satisfies $v(\mu'_i) = v(\mu_i)$; the latter is a consequence of conditionally strict preferences. As $v^{\mathcal{P}}$ is linear in probabilities and \tilde{v} is a maximum over all partitions, it is easy to see that $\alpha\tilde{v}(\mu') + (1 - \alpha)\tilde{v}(\mu'') > \tilde{v}(\mu)$, where $\alpha \in (0, 1)$ satisfies $\alpha\mu' + (1 - \alpha)\mu'' = \mu$. This geometry has a familiar informational interpretation: starting from any $\lambda^{\mathcal{P}}$, the sender can almost always 'split' posteriors in such a way that the receiver's induced action either remains unchanged, or switches to an action the sender prefers.²⁷

In reality, partitional communication does appear to be common. For instance, ratings agencies rank debt instruments by creditworthiness, schools grade students, and health agencies categorize restaurants by hygiene. Moreover, the use of these categories

²⁵Such payoffs are clearly generic, in the sense that they form an open and dense set in $\mathbb{R}^{|\Theta| \times |A|}$.

²⁶Lipnowski and Ravid (2020) (Corollary 3) identify a related result for one-shot cheap talk.

²⁷The class of problems studied in Proposition 3 is therefore one in which solving the commitment problem can be hard (Lipnowski and Mathevet (2017b)) and yet evaluating the geometry of \tilde{v} simple.

often appears relatively stable even as the underlying environment changes. For example, investment grade bonds maintain low default rates even during recessions.²⁸ When can such categorization attain Bayesian persuasion payoffs, irrespective of prior beliefs?

Proposition 4 *A partition $\mathcal{P} \in \mathcal{P}^\Theta$ satisfies $v^\mathcal{P} = \tilde{v} = V_{kg}$ if and only if:*

(i) $v^\mathcal{P} \geq v$, and

(ii) v is concave on ΔP_i , $i = 1, \dots, k$.

Proposition 4 provides necessary and sufficient conditions for communication over a single partition to achieve the Bayesian persuasion payoff irrespective of priors. Of course, these conditions are also sufficient for \tilde{v} to be concave. Condition (i) compares \mathcal{P} with a single alternative partition: it requires that the sender always (weakly) prefer communicating partition \mathcal{P} to nothing at all. By contrast, it imposes nothing on the relative payoff from other partitions. The second condition requires that the sender's payoff be concave on the set of beliefs that correspond to a receiver learning $\{\theta \in P_i\}$, for each element of the partition. An application of Jensen's inequality shows that under condition (ii), $v^\mathcal{P}$ is concave on $\Delta\Theta$. But then, by an almost identical argument to that in Theorem 1, condition (i) implies that $v^\mathcal{P}$ and V_{kg} must be the same.

These conditions can be interpreted in terms of the nature of conflicts of interest between the sender and receiver. Condition (i) represents an alignment of interests between sender and receiver: they both agree it is always better to share information about the element $P_i \in \mathcal{P}$ in which θ lies. However, their alignment ends here. While a receiver wants ever more information, condition (ii) implies the sender never wants to provide any further information distinguishing realizations of the state within P_i .

2.3 Imperfect Monitoring

The results of Fudenberg and Levine (1994) imply that under imperfect monitoring the costs of long-run persuasion only increase. Indeed, even partitional information structures may now attract on-path punishments. We use example 1 to illustrate how

²⁸To the extent that recessions represent a fall in the prior 'quality' of bonds, this stability suggests that 'investment' grades are categorized according to the same fundamentals throughout.

these results apply in our setting. Consider an equilibrium in which the sender provides full information whenever possible (i.e., when not being punished). To prevent mis-selling (announcing ‘like new’ when $\theta = l$), the seller must be punished for a bad review ($\omega = b$). Since this would occur with probability p , the minimal punishment cost κ needed to support full information must satisfy the indifference condition $1 = \delta p \kappa$. However, unlike perfect monitoring, the seller now also suffers a bad review when $\theta = h$ and thus pays κ with probability $1 - p$. As this is true at *every* history where his advice is credible, his equilibrium payoff is bounded above by $\mu_0 - \mu_0 \frac{1-p}{p}$. Moreover, as punishments necessitate less information, imperfect monitoring harms customers too; one can show their (discounted) payoffs are also bounded above by $\mu_0 - \mu_0 \frac{1-p}{p}$.

3 Persuasion with Incomplete Public Records

It is natural to conjecture that by providing receivers less information on which to base punishments, incomplete public records only exacerbate the problems identified in sections 2.2 and 2.3. Indeed, in the limiting case of the least informative record only stage Nash equilibria can be sustained. Yet, in many markets coarse summaries such as star ratings, badges, or certificates are adopted over more informative alternatives. In this section, we introduce a third-party designer (the ‘platform’) who can commit to rules that determine what incoming receivers observe about \underline{h}_t . With such a designer, we show how any information structure that the sender prefers to one-shot cheap talk can be implemented with almost no cost. But first, we illustrate how pooling histories can increase equilibrium payoffs by persuading receivers to adopt strategies that would not be incentive compatible could they observe the complete feedback history.

3.1 Pooling Histories to Persuade

In the context of example 1 with imperfect monitoring, we show that pooling feedback histories for a single receiver can strictly increase average equilibrium payoffs. In particular, the customer at $t = 2$ observes no information about $t = 1$, but all other customers observe \underline{h}_t . That is, $r_2(\underline{h}_2) = \emptyset$, $\underline{h}_2 \in \underline{\mathcal{H}}_2$, and $r_t(\underline{h}_t) = \underline{h}_t$ otherwise. As no other part of the game is changed, any increase in payoffs must arise because the coarse record ‘persuades’ the $t = 2$ customer to adopt a behavior incompatible with her observing \underline{h}_2 .

Proposition 5 *If receivers observe an incomplete public record, there is a PBE with payoffs of sender and receiver exceeding the bounds of Fudenberg and Levine (1994).*

In Appendix B, we construct a PBE in *private strategies* with average seller and customer payoffs exceeding the uniquely efficient ‘full information’ equilibrium from section 2.3. The seller provides full information at $t = 1$ and again at $t = 2$ if $\omega_1 = b$; yet, if $\omega_1 \neq b$, he sends $m_2 = \text{‘like new’}$ irrespective of θ_2 . In each of these periods, the customers buy if and only if $m_t = \text{‘like new’}$. From $t = 3$, play is determined as follows: If $\omega_1 \neq b$, players progress to the full information continuation of section 2.3, regardless of ω_2 . However, if $\omega_1 = b$ the seller is punished with a babbling equilibrium for a fixed number of periods before reverting to ‘full information’; if $\omega_2 = b$ as well, then this punishment phase is lengthened. For a range of parameters, this is an equilibrium. In particular, since the customer at $t = 2$ does not observe ω_1 and the seller tells the truth when $\omega_1 = b$, she finds a ‘like new’ message credible enough to buy if $\omega_1 = b$ is likely and μ_0 is not too low. In this way, credibility is extended—punishment free—from histories at which the seller is incentivized ($\omega_1 = b$) to ones where he is not ($\omega_1 \neq b$).

Relative to the full information outcome of section 2.3, this equilibrium guarantees the seller a punishment-free sale at $t = 2$ if he avoids a bad review in period 1. Necessarily, this continuation payoff exceeds the Fudenberg and Levine (1994) bound for the seller. Moreover, this higher continuation also provides an extra incentive to avoid a bad review at $t = 1$, reducing the required punishment associated with credible communication at $t = 1$. Finally, less punishment (less babbling) also benefits customers. We show this can more than compensate for the cost of occasionally buying from an uninformative seller—increasing average customer payoffs too.

These gains relied on two features of the incomplete public record: the pooling of the feedback histories $\omega_1 = b$ and $\omega_1 \neq b$ at $t = 2$, and their subsequent separation for $t > 2$. The pooling implies that the private strategy induces a single ‘average’ distribution over posteriors (call it λ_2) at both $\omega_1 \neq b$ and $\omega_1 = b$, persuading the customer to adopt a strategy that would be suboptimal could she observe h_2 . The subsequent separation allows the seller’s behavior to be incentivized at the two histories as if the record were complete. In this manner, λ_2 can be induced via two pure stage strategies, each cheaper

to incentivize than a public strategy inducing λ_2 .^{29,30}

More generally, any incomplete record that improves payoffs must persuade some receiver to adopt a behavior that would violate incentive compatibility could she observe the complete feedback history \underline{h}_t . To see this, consider any incomplete record with a corresponding equilibrium strategy profile that yields payoffs $(u, v) \notin \mathcal{E}(\mu_0, \delta)$. Clearly, there is a strategy profile when $r_t = \underline{h}_t$ that induces exactly the same distribution over messages and actions at every \underline{h}_t . By construction, this profile cannot be an equilibrium with complete records, i.e., it must violate incentive compatibility for someone. As the sender’s incentives are identical in both cases (given receivers’ strategies), some receiver must have a profitable deviation. Hence, any improvement must rely on altering the receivers’ incentives by controlling the information they have about history.

3.2 The Commitment Set by Design

Increasingly, online platforms use badge systems to publicly record participants’ previous interactions. These systems monitor the participant during *evaluation phases*, collecting reviews and complaints about misleading descriptions. After each evaluation, the platforms publicly award or withdraw a ‘badge’ based on whether the distribution of reviews or complaints meets a set of *standards* decided by the platform. Finally, they hide some feedback data from the customers, while making it available to the suppliers.

Consider a platform which can design the public record, $(r_t, \mathcal{R}_t)_t$. A **badge system** is a *binary* public record $(r_t, \mathcal{R}_t)_t$ where $\mathcal{R}_t = \{\mathcal{G}, \mathcal{B}\}$, $\forall t$. We interpret \mathcal{G} as a public badge given to the sender (he is a ‘badge bearer’) and \mathcal{B} as absence thereof. A badge system is **simple** if it can be described by three parameters: (i) the length of **evaluation phases**, $\Gamma \in \mathbb{N}$; (ii) the length of **suspension phases**, $\beta \times \Gamma \in \mathbb{N}$; and (iii) a **set of standards**, \mathcal{S} , that reflect the outcomes required for a sender to retain a badge. More precisely, the standards define, for each possible $m \in \Delta\Theta$ and all $\omega \in \Omega$, a set $\mathcal{S}(m)$ of

²⁹By contrast, if ω_1 remained hidden in all subsequent periods, the seller’s continuation payoff from any strategy would be independent of ω_1 . In this case, for any equilibrium in private strategies there is a payoff equivalent one in public strategies, and the bounds of section 2 must apply. This is the same reason that private strategies cannot increase payoffs when the record is complete. See Lemma 2 in Appendix A for the argument with complete public records.

³⁰This also implies that incomplete records can improve payoffs only if they support new equilibria in private strategies. Indeed, any PPBE with an incomplete record is also a PPBE when it is complete.

allowable joint frequencies $s(m)$ over (m, ω) .³¹ As each message induces an action in equilibrium, these can be interpreted as targets for receivers' reported experiences.

Given these parameters, a **simple badge system** works as follows: at $t = 1$, the sender is given the badge ($r_1 = \mathcal{G}$). This initiates an evaluation phase, throughout which the badge is retained. At the end of an evaluation phase, feedback obtained during that phase is compared against the standards. Specifically, let

$$\ell_t(m, \omega) = \frac{1}{\Gamma} \sum_{\tau=t-\Gamma+1}^t \mathbb{1}((m_\tau, \omega_\tau) = (m, \omega)) \quad (3)$$

be the realized frequency of (m, ω) in the Γ interactions preceding $t \geq \Gamma$, where $\mathbb{1}$ is the indicator function, and let $\ell_t(m) = (\ell_t(m, \omega))_{\omega \in \Omega}$. Then if $\ell_\Gamma(m) \notin \mathcal{S}(m)$ for some m , the sender loses the badge for $\beta \times \Gamma$ periods. Otherwise, his badge is renewed. After this, the entire process restarts.

We aim to understand the scope for these systems to expand the set of payoffs achievable with communication. To this end, we consider a setting in which arriving receivers do not initially know the date on which the sender is next evaluated. Similar to Kremer et al. (2014), we assume receivers arrive uniformly at random during evaluation phases and do not observe t . This simplifies analysis and only further limits receivers' information about \underline{h}_t , thereby leveraging the insight of section 3.1. In section 4, we show how platforms can create similar uncertainty if arrivals are nonrandom. We fix $\mu_0 \in \text{int}\Delta\Theta$, and make two additional assumptions:³²

Assumption 2 *The distributions $(p(\theta \mid \omega))_{\theta \in \Theta}$ are linearly independent across Ω .*

Assumption 2 slightly relaxes the standard identifiability condition in repeated games (recall we allow some $a' \in A$ to be uninformative about θ). If this assumption is violated, the sender's average strategy cannot be accurately monitored even with infinite data. Denoting the closure of a set X by $\text{cl}(X)$, we also assume the following:

Assumption 3 *For any $\mu \in \Delta\Theta$, there exists an open set $X_\mu \subset \Delta\Theta$ such that $\mu \in \text{cl}(X_\mu)$, and $a(\mu)$ (and hence $v(\mu)$) are continuous functions on $\text{cl}(X_\mu)$.*

³¹Each $s(m)$ is a feasible vector of joint probabilities over (m, ω) , $\omega \in \Omega$, whose sum must of course not exceed 1. Formally each $s(m) \in \Delta_c \Omega := \{s \in \mathbb{R}^{|\Omega|} : \sum_{\omega \in \Omega} s(m, \omega) \leq 1, s \geq 0\}$.

³² $\mu_0 \in \text{int}\Delta\Theta$ is without loss of generality for payoffs: simply exclude any θ for which $\mu_{0,\theta} = 0$. This ensures the continuity properties of $v(\mu)$ extend to $v(\lambda)$ on $\Lambda(\mu_0)$ appropriately (see Lemma 14).

Assumption 3 is a mild restriction. It allows for discontinuities in $v(\mu)$, such as in example 1, but rules out knife-edge cases where the receiver’s action (and hence v) changes radically after *every* small perturbation.³³

Under these assumptions, we establish the value of simple badge systems. Let $\underline{\Lambda}(\mu_0)$ be the subset of $\Lambda(\mu_0)$ for which $v(\lambda)$ exceeds the sender’s worst one-shot cheap talk payoff, and let $\underline{\mathcal{CS}} \subset \mathcal{CS}$ be the associated payoffs:

Theorem 2 *Let Assumptions 2 - 3 hold, and $\mu_0 \in \text{int } \Delta\Theta$. For any $(u', v') \in \underline{\mathcal{CS}}$, there is a $\lambda^* \in \underline{\Lambda}(\mu_0)$ and a limiting PBE with a simple badge system in which, as $\delta \rightarrow 1$,*

1. *in an evaluation phase, a badge bearer’s strategy implies a time-averaged information structure*

$$\lambda^* = \frac{1}{\Gamma} \sum_{\tau=1}^{\Gamma} \mathbb{E}[\lambda_{\tau} | \mathcal{G}],$$

2. *average payoffs are $(u(\lambda^*), v(\lambda^*)) = (u', v')$.*

Theorem 2 shows that the benefits of coarse rating systems like those found online can be large: they can efficiently solve imperfect monitoring problems, driving on-path punishment close to zero. In stark contrast to section 2, simple badge systems allow a sender to capture the full value of his communication; indeed, he achieves the payoff $v(\lambda^*)$ with a strategy that induces a time-averaged information structure of λ^* . Of course, Theorem 2 implies badge systems can be used to implement Bayesian persuasion payoffs, but their scope is much greater—they can support *any* payoff profile in \mathcal{CS} that the sender prefers to one-shot cheap talk. For instance, in the patient limit, every feasible payoff in example 1 (Figure 1b) can be supported with some set of standards.

As $(u', v') \in \mathcal{CS}$, there is a finite-support λ^* with $(u(\lambda^*), v(\lambda^*)) = (u', v')$. In the proof, we construct a corresponding sequence of simple badge systems that admit ‘good’ equilibria as $\delta \rightarrow 1$. These systems involve long evaluation phases and standards which allow only distributions of (m, ω) close to the one induced by λ^* . Under Assumption 2, we can choose standards tough enough that if the sender uses a time-average information structure that differs more than a little from λ^* , he loses his badge, and yet lenient enough that if he adopts λ^* each period, he keeps his badge with probability close to 1.

³³In writing $a(\mu)$, we implicitly assume the receiver’s best response is unique on $cl(X_{\mu})$. Its continuity on that set shortens the proof of Theorem 2 by ensuring the existence of equilibria where receiver strategies are pure. The continuity of $v(\mu)$ is necessary for our argument.

Of course, in *any* equilibrium, suspensions induce one-shot cheap talk incentives. To analyze a badge bearer’s behavior, we first adapt an argument from Radner (1985): if receivers’ beliefs satisfy $\mu(m) \approx m$, suspensions are long enough, and he is patient, then his optimal strategy is to play close to λ^* on average to avoid the punishment of a suspension. We then consider a receiver’s beliefs upon observing a message m . The only historical data she can condition on is $r = \mathcal{G}$, so her beliefs can depend only on the sender’s average strategy. Hence, as this is close to λ^* , her beliefs must indeed satisfy $m \approx \mu$. Building on this intuition, the proof shows best responses and beliefs can be mutually closed in arbitrarily small intervals of each other. Along with Assumption 3, we then prove existence via a fixed-point argument.

While it is well known that it can be more efficient to discipline an agent’s average behavior—for instance, via review strategies (Radner, 1985) or linking decision problems (Jackson and Sonnenschein, 2007; Escobar and Toikka, 2013)—such approaches are not possible when the record is complete. This is because these strategies require receivers take actions that are not incentive compatible at certain feedback histories. For example, the lenient standards of the SBS imply there must be some histories at which the sender can send any message without fear of punishment. If the receiver knew she were at such a history, she would deviate (see Example 2). The SBS resolves this problem by hiding these histories from the receivers. As we argued in section 3.1, this kind of persuasion is necessary for any incomplete record to support payoffs exceeding the Fudenberg and Levine (1994) bound. Theorem 2 shows these gains can be large.

We illustrate Theorem 2 with a simple badge system designed to attain the Bayesian persuasion payoff for the seller in the online market setting. In what follows, we focus on behavior during evaluations (recall suspensions induce babbling).

Example 2 Consider example 1 with $\mu_0 = \frac{1}{3}$, $\underline{\mu} = 0.5$, and arbitrary $p > 0.5$. Let $\epsilon = \Phi(-6) \cong 0.00034$, where Φ is the standard normal c.d.f. Consider a simple badge system with standards that impose an upper limit on the frequency of bad reviews: the seller keeps his badge if $\ell_\Gamma(\text{‘like new’}, b) \leq \bar{\ell} := \frac{1}{3} - \frac{1}{6}\epsilon p$. Set $\beta = 1$ and let Γ be some large integer satisfying $6\sqrt{\frac{\hat{\ell}(1-\hat{\ell})}{\Gamma}} \leq \frac{1}{6}\epsilon p$, where $\hat{\ell} = \frac{1}{3}$ is the expected frequency of feedback (‘like new’, b) under the stage strategy associated with λ^{kg} .³⁴ This system admits a limiting equilibrium with an average payoff within ϵ of $V_{kg}(\mu_0) = \lambda^{kg}(\underline{\mu}) = \frac{2}{3}$.

³⁴For any λ' with support $\{0, \mu'\}$, the probability of feedback (μ', b) is $\mu_0(1 - 2p) + \lambda'(\mu')p$.

To illustrate the argument, we examine the seller's incentives over a simplified class of strategies that send 'like new' and 'used' with constant conditional probabilities throughout the evaluation phase. To this end, let σ^* be a constant strategy analogous to choosing λ^* every period, where $\text{supp } \lambda^* = \{0, \mu^*\}$, $\mu^* = \underline{\mu}(1 + \frac{\epsilon}{2-\epsilon})$, and $\lambda^*(\mu^*) = \lambda^{kg}(\underline{\mu})(1 - 0.5\epsilon)$. Suppose the seller adopts σ^* and that $\mu(\text{'like new'}) \geq 0.5$. As Γ is large, the central limit theorem implies the distribution of $\ell_\Gamma(\text{'like new'}, b)$ is approximately $\mathcal{N}(\ell^*, \frac{\ell^*(1-\ell^*)}{\Gamma})$, where $\ell^* = \bar{\ell} - \frac{1}{6}\epsilon p$ is more than 6 standard deviations below $\bar{\ell}$. Thus, the seller keeps his badge with probability close to $1 - \epsilon$. It is then easy to show that, as $\delta \rightarrow 1$, the seller can secure an average payoff of at least $\frac{2}{3} - \epsilon$ with σ^* .

Consider now a deviation to the constant strategy $\hat{\sigma}$ associated with λ^{kg} . During evaluation, the increase in sales from this—or indeed any—deviation must be bounded above by $(1 - \lambda^*(\mu^*))\Gamma \approx \frac{1}{3}\Gamma$ (the increase from sending 'like new' with probability 1 every period). Yet, as $\hat{\ell}$ and ℓ^* are symmetric about $\bar{\ell}$, the deviation increases the probability of a suspension by at least $1 - 2\epsilon \approx 1$, costing the seller an expected $\lambda^*(\mu^*)\Gamma \approx \frac{2}{3}\Gamma$ sales. Thus, as $\delta \rightarrow 1$, deviating to $\hat{\sigma}$ cannot be profitable. This applies a fortiori to strategies that target a rate of sales greater than $\lambda^{kg}(\underline{\mu})$, because the seller loses the badge with an even higher probability. Thus, with the restriction to constant strategies, the supposition of $\mu(\text{'like new'}) \geq 0.5$ is justified.

Of course, a constant strategy cannot be optimal. To see why, suppose in the final period of an evaluation phase that $\ell_{\Gamma-1}(\text{'like new'}, g) > \ell^*$. As the seller will keep his badge regardless of his final report, he will claim 'like new' irrespective of θ . Hence, the seller will adopt a private feedback-dependent strategy. Nonetheless, a similar argument can be applied. We show that to meet the standards with high probability, he must adopt a strategy that is close to σ^* 'on average'. That is, averaging the strategy over time, the ex-ante probability of sending message m given θ is close to $\sigma^*(m | \theta)$. Hence, because the badge system keeps customers in the dark about both arrival time and feedback history, customers form beliefs as if they were facing a constant strategy close to σ^* .

A similar argument to that in example 2 applies to any payoff profile in \mathcal{CS} . Clearly though, each such profile is enforced by a different set of standards. For instance, if the platform wished to implement full information payoffs then it would have to replace the threshold $\bar{\ell}$ with $\frac{1}{3}(1 - p)$; the choice of standards affects the distribution of surplus between market participants.³⁵ This suggests that in reality, platforms' ratings

³⁵Any payoff on the efficient frontier of \mathcal{CS} can be sustained with an upper-threshold on 'bad'

systems—and the resulting communication on them—may be shaped by the platforms’ relative incentives to attract sellers and customers.

Finally, we briefly discuss two aspects of our analysis. First, whether complete or incomplete, public record systems do require some source of commitment themselves. It is not our aim to explain that commitment but to tease out its implications for an uncommitted sender. Nonetheless, we note that the systems discussed are partitions of history. Thus, long-run incentives may serve as an efficient source of commitment. Moreover, the platform faces many copies of the same problem; hence, similar to Jackson and Sonnenschein (2007), the distribution of badges could act as the source of discipline. Second, unlike the mediators typically analyzed in repeated games, our platform cannot commit to private messages. This reflects most online platforms, who take a hands-off approach to individual transactions, perhaps because mediating every single interaction is too costly. By ignoring mediation, we might worry that something is lost. Yet despite limiting the platform’s power, Theorem 2 shows that badge systems can still implement all commitment payoff profiles preferred to one-shot cheap talk.

4 Discussion: Alternative Review Systems

In section 3.2 the platform’s control over public records was strong: it could privately collect feedback from receivers, it could prevent receivers from accessing this feedback directly, and it benefited from the random arrival of receivers (further limiting their information). Moreover, communication was the only source of moral hazard. Here we discuss how review systems with less control can still improve outcomes. In doing so we also highlight the costs and benefits of each one relative to simple badge systems.

4.1 Deterministic Arrival Systems

In section 3.2 receivers could not infer the time until the next evaluation. While information on evaluation dates is often searchable, customer behavior online seems better approximated by uncertainty over such details. Our results support their apparent disinterest: eBay’s strict standards incentivize badged sellers to almost always behave.

feedback, but we need the tight two-sided standards to get the whole set (See proof of Theorem 2).

Hence even small attention costs would deter a search. Yet with more relaxed standards, incentives to find out these dates could rise. We briefly discuss the implications of this here. The online appendix provides formal results in the setting of example 1.

As one might expect, such information can constrain the performance of some simple badge systems. Nonetheless, when t is observable, the platform can replicate the effect of random arrivals with a system that instead randomizes evaluation dates directly. If the platform hides the (randomly chosen) evaluation dates from customers, it is *as if* they had arrived randomly (see Proposition 7). Even if customers *must* be informed of evaluation dates, the platform can still replicate the required uncertainty by adopting more complex standards that depend on *full sequences* of outcomes rather than simple averages (see Proposition 8).³⁶ The idea is to make the sender endogenously create the uncertainty by mixing between two non-stationary strategies.³⁷ As with the SBS, such systems still require receivers not observe the complete record to improve payoffs.

We view stochastic evaluation dates as just another interpretation of our random arrival assumption. However, simple badge systems have three advantages over more complex evaluations. First, they are easier to identify and implement. Second, they do not rely on common knowledge that the sender mixes with precisely the right probabilities. Finally, simple badge systems do not require exact knowledge of δ .

4.2 Blind Sender Review Systems

In many cases the platform’s commitment to badge systems is easy to motivate. But in the absence of an ability to discipline a third-party platform, customers could instead maintain a distributed ledger of feedback themselves. Hence, we now consider whether payoffs can be improved when receivers observe all feedback but the sender only has limited access. We call this a **blind sender** review system. At each t the incoming customer observes \underline{h}_t . Badges are again awarded to the seller based on evaluations of standards, but now by the customers rather than a third party. The seller observes the badges provided but not the individual feedback, ω_t . Notice, such systems fall outside our definition of public records; as they hide information from the sender rather than

³⁶We thank an anonymous referee for pointing us towards this system, which draws nicely on the linking decisions insight of Jackson and Sonnenschein (2007) and Escobar and Toikka (2013).

³⁷Our construction is reminiscent of a mediated equilibrium proposed in Sugaya and Wolitzky (2018), with the additional constraint of sender indifference across strategies.

the receivers, they improve payoffs via a different channel.

Proposition 6 *Consider example 1 with imperfect monitoring, $p < 1$. Using a blind sender review system, full information payoffs (μ_0, μ_0) are attainable in the limit as $\delta \rightarrow 1$. The sender cannot earn payoffs greater than μ_0 in any public PBE.*

The proof uses a standard in which the seller loses his badge if and only if *every* product sold receives bad feedback. Building on Abreu et al. (1991) and Fuchs (2007), we construct an equilibrium in which a badged seller is *continually* disciplined to be honest by ‘reusing’ the *same* threat of punishment, dramatically reducing their aggregate cost. Moreover, there is now clearly no need to hide history from receivers.

Blind sender systems have two major limitations. As Proposition 6 shows, reusing punishments alone cannot costlessly support non-partitional information structures. When customers observe a finer feedback history than the seller, the logic of Proposition 2 still applies. Moreover, they can be fragile: if a seller can find out that he has even a single good review, he can lie for the remainder of an evaluation without cost. This may be one reason why sellers do observe the history of feedback on most platforms.

4.3 Quota-Based Review Systems

Some platforms may receive very little feedback. Even without feedback, Margaria and Smolin (2018) show how to obtain a folk theorem in a setting with a long-run receiver and a sender with state-independent preferences. In the context of example 1, their equilibrium constructions make the seller indifferent between all *sequences* of messages over long horizons. To generate indifference, the customer is required to occasionally buy regardless of the seller’s claims. This is obviously not feasible with short-run customers and complete public records. However, as the frequency of such ‘bad purchases’ is low as $\delta \rightarrow 1$, we conjecture a badge system may be able to achieve any payoffs in \mathcal{CS} . Even if this is true, feedback-free ratings would still have limitations: they require exact knowledge of δ (to generate the indifference) and are not robust to state-dependent payoffs. By contrast, simple badge systems are robust to both issues.

4.4 Simple Badge Systems for Other Forms of Moral Hazard

Online platforms face many forms of moral hazard beyond communication: sellers may fail to deliver on time, hosts may fail to keep accommodation clean, or workers may shirk. Simple badge systems can equally solve these problems. For the platform, the key difference between communication and these classic moral hazards lies in how much players can infer beyond the badge. In our setting additional historical information may leak out (receivers can update after seeing the sender’s message). But classic moral hazards do not suffer this leakage, so the argument underpinning Theorem 2 applies to them *a fortiori* (see Corollary 1 in the online appendix). Indeed, subject to adjustments of Assumptions 2 and 3, simple badge systems ought to support commitment payoffs for any two-player stage game in the setting of Fudenberg and Levine (1994), as above.

5 Related Literature

Since Kamenica and Gentzkow (2011), a large body of work has sought to understand strategic communication when a sender has exogenous commitment power.³⁸ We contribute to this literature by relaxing the commitment assumption. Perez-Richet (2014), Salamanca (2021), Lipnowski et al. (2019), and Perez-Richet et al. (2021) study persuasion when the sender has partial commitment power. Mathevet et al. (2019) consider how well repetition substitutes for commitment when the sender may be an exogenous ‘commitment type’. We examine persuasion when there is no hope of the sender having exogenous commitment power. Moreover, we show that the design of coarse public rating systems (and repetition) can substitute for a complete lack of commitment on the sender’s part. Like us, Chakraborty and Harbaugh (2010) and Lipnowski and Ravid (2020) analyze persuasion when the sender cannot commit. Relative to them, our contribution is to introduce an endogenous source of commitment.

Partitional information structures have gained attention for their simplicity, applicability and tractability. Kolotilin (2018), Dworzak and Martini (2019), Mensch (2021), and Kolotilin and Li (2019) identify conditions under which *monotone* partitional in-

³⁸For instance, Rayo and Segal (2010), Taneva (2019), Bergemann and Morris (2013), Ely (2017), and Bizzotto et al. (2021).

formation structures are optimal.³⁹ We identify a new, enforcement-based appeal of partitional information structures and provide new conditions for their optimality.

We also relate to recent work on cheap talk in repeated games. Hörner et al. (2015) and Margaria and Smolin (2018) examine such games between long-run players, and obtain a folk theorem. Like us, Jullien and Park (2020) examine a setting with short-run receivers. These papers all consider settings where truthful equilibria are without loss. In our setting there is a meaningful role for persuasion. Aumann and Maschler (1995) study a repeated zero-sum game with two infinitely patient agents and asymmetric information about *perfectly persistent* stage payoffs, showing Bayesian persuasion is a credible cheap talk outcome. We focus on complementary settings where the state is changing, receivers are short-lived, and $\delta < 1$. Kuvalekar et al. (2021) study a problem similar to ours, but with no monitoring at all of the sender’s accuracy.

Other papers have examined hiding history in repeated games. Sugaya and Wolitzky (2018) shows how private monitoring with a mediator can help sustain a collusive agreement by making it more difficult for firms to tailor deviations to market conditions in a stochastic game. To isolate this channel, they assume perfect monitoring. By contrast, we show how to eliminate inefficiencies caused by imperfect monitoring in stationary games.⁴⁰ Abreu et al. (1991) and Fuchs (2007) hide history from *long-run* players when monitoring is imperfect. Unlike them, we cannot exploit ‘reusable’ punishments (see section 4.2). Instead, we show how receivers can be persuaded to provide more efficient forms of incentive, allowing us to tackle *both* the imperfect monitoring and mixing problems present in Fudenberg and Levine (1994). Bhaskar and Thomas (2019) provide a coarsening of records that mitigates inefficiencies in a repeated investment game caused by bounded memory. While their inefficiencies could also be solved by providing complete public records, our SBS can strictly improve on complete records.

Several papers, going back to Myerson (1986), have designed information in extensive-form games to improve payoffs. Recently, Gershkov and Szentes (2009) apply this to voting mechanisms, and Kremer et al. (2014), to social learning. We show how this idea can be used to improve persuasion and the provision of efficient incentives in repeated

³⁹By contrast, non-partitional information appears in the Bayesian persuasion solutions to example 1, Brocas and Carrillo (2007), Gill and SgROI (2008, 2012), Rayo and Segal (2010) and Alonso and Câmara (2016).

⁴⁰While their main results apply to a stochastic game, they do provide a numerical example where mediation can improve payoffs in a stationary environment. In this example, monitoring is perfect.

games more generally. Doval and Ely (2020) provide related tools for finite games, but these do not apply to our infinite-horizon setting.

Dellarocas (2005) and Tadelis (2016) examine how online review systems can implement the Fudenberg and Levine (1994) bound. Our results show how review systems can do better. Ekmekci (2011) shows how reducing a rating system’s memory can prevent the collapse of reputations, à la Cripps et al. (2004), and improve seller payoffs; but the existence of reputational types is necessary for making improvements. By contrast, our badge systems yield Pareto improvements without any types at all.

6 Conclusion

We have analyzed a model with a meaningful role for persuasion where commitment lies in repetition and a public record of past accuracy. When the public record is complete, we gave necessary and sufficient conditions on the underlying persuasion problem under which a loss of ex ante commitment is costless. Unfortunately, even if the sender’s private information can be observed ex post, these conditions are often not met. However, by applying information design tools to the public record, we can costlessly support any form of communication the sender prefers to one-shot cheap talk.

We conclude with some open questions. First, instead of examining cheap talk, we could have examined long-run persuasion when the sender’s private information is endogenously acquired, partially verifiable, or both. It is possible to construct examples where such changes significantly alter communication outcomes to the detriment of both parties. Second, how can commitment to the rules that generate the public record be justified? As is common in the literature, this paper treated such commitment as given. However, we suggested two routes toward justification: the platform as a third, long-run player and the possibility of distributed ledgers (section 4.2). With the rise of blockchain technologies, exploring the latter would also be interesting.

References

Dilip Abreu, David Pearce, Ennio Stacchetti, et al. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58(5):1041–1063, 1990.

- Dilip Abreu, Paul Milgrom, and David Pearce. Information and timing in repeated partnerships. *Econometrica*, 59(6):1713–1733, 1991.
- Ricardo Alonso and Odilon Câmara. Persuading voters. *The American Economic Review*, 106(11):3590–3605, 2016.
- Robert J Aumann and Michael B Maschler. *Repeated games with incomplete information*. MIT press, 1995.
- Heinz Bauer. *Probability Theory*. New York: De Gruyter, 2011.
- Dirk Bergemann and Stephen Morris. Robust predictions in games with incomplete information. *Econometrica*, 81(4):1251–1308, 2013.
- James Best and Daniel Quigley. Persuasion for the long run. *Working Paper*, 2017.
- V Bhaskar and Caroline Thomas. Community enforcement of trust with bounded memory. *The Review of Economic Studies*, 86(3):1010–1032, 2019.
- Patrick Billingsley. *Probability and measure*. John Wiley & Sons, 3. ed edition, 1995.
- Jacopo Bizzotto, Jesper Rüdiger, and Adrien Vigier. Dynamic persuasion with outside information. *American Economic Journal: Microeconomics*, 13(1):179–94, 2021.
- Isabelle Brocas and Juan D Carrillo. Influence through ignorance. *RAND Journal of Economics*, 38(4):931–947, 2007.
- Archishman Chakraborty and Rick Harbaugh. Persuasion by cheap talk. *American Economic Review*, 100(5):2361–2382, 2010.
- Martin W Cripps, George J Mailath, and Larry Samuelson. Imperfect monitoring and impermanent reputations. *Econometrica*, 72(2):407–432, 2004.
- Gerard Debreu. A social equilibrium existence theorem. *Proceedings of the National Academy of Sciences*, 38(10):886–893, 1952.
- Chrysanthos Dellarocas. Reputation mechanism design in online trading environments with pure moral hazard. *Information Systems Research*, 16(2):209–230, 2005.
- Laura Doval and Jeffrey C Ely. Sequential information design. *Econometrica*, 88(6):2575–2608, 2020.

- Piotr Dworzak and Giorgio Martini. The simple economics of optimal persuasion. *Journal of Political Economy*, 127(5):1993–2048, 2019.
- Mehmet Ekmekci. Sustainable reputations with rating systems. *Journal of Economic Theory*, 146(2):479–503, 2011.
- Jeffrey C Ely. Beeps. *The American Economic Review*, 107(1):31–53, 2017.
- Juan F Escobar and Juuso Toikka. Efficiency in games with markovian private information. *Econometrica*, 81(5):1887–1934, 2013.
- Ky Fan. Fixed-point and minimax theorems in locally convex topological linear spaces. *Proceedings of the National Academy of Sciences of the United States of America*, 38(2):121, 1952.
- William Fuchs. Contracting with repeated moral hazard and private evaluations. *American Economic Review*, 97(4):1432–1448, 2007.
- D Fudenberg and D Levine. Efficiency and observability with long run and short run players. *Journal of Economic Theory*, 62:103–135, 1994.
- Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, Cambridge, MA, 1991.
- Drew Fudenberg, David M Kreps, and Eric S Maskin. Repeated games with long-run and short-run players. *The Review of Economic Studies*, 57(4):555–573, 1990.
- Drew Fudenberg, David Levine, and Eric Maskin. The folk theorem with imperfect public information. *Econometrica: Journal of the Econometric Society*, pages 997–1039, 1994.
- Alex Gershkov and Balázs Szentes. Optimal voting schemes with costly information acquisition. *Journal of Economic Theory*, 144(1):36–68, 2009.
- David Gill and Daniel SgROI. Sequential decisions with tests. *Games and Economic Behavior*, 63(2):663–678, 2008.
- David Gill and Daniel SgROI. The optimal choice of pre-launch reviewer. *Journal of Economic Theory*, 147(3):1247–1260, 2012.
- Irving L. Glicksberg. A further generalization of the kakutani fixed point theorem, with application to nash equilibrium points. *Proceedings of the American Mathematical Society*, 3(1):170–174, 1952.

- Johannes Hörner, Satoru Takahashi, and Nicolas Vieille. Truthful equilibria in dynamic bayesian games. *Econometrica*, 83(5):1795–1848, 2015.
- Xiang Hui, Maryam Saeedi, Zeqian Shen, and Neel Sundaresan. Reputation and regulations: evidence from ebay. *Management Science*, 62(12):3604–3616, 2016.
- Matthew O Jackson and Hugo F Sonnenschein. Overcoming incentive constraints by linking decisions 1. *Econometrica*, 75(1):241–257, 2007.
- Bruno Jullien and In-Uck Park. Communication, feedbacks and repeated moral hazard with short-lived buyers. *Working Paper, Available at SSRN 3655043*, 2020.
- Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.
- Anton Kolotilin. Optimal information disclosure: A linear programming approach. *Theoretical Economics*, 13(2):607–635, 2018.
- Anton Kolotilin and Hongyi Li. Relational communication. *Working Paper*, 2019.
- Ilan Kremer, Yishay Mansour, and Motty Perry. Implementing the Wisdom of the Crowd. *Journal of Political Economy*, 122(5):988–1012, 2014.
- Aditya Kuvalekar, Elliot Lipnowski, and Joao Ramos. Goodwill in communication. *USC Marshall School of Business Research Paper Sponsored by iORB, No. Forthcoming*, 2021.
- Elliot Lipnowski and Laurent Mathevet. Simplifying bayesian persuasion. *Unpublished Paper, Columbia University.[642]*, 2017a.
- Elliot Lipnowski and Laurent Mathevet. Simplifying bayesian persuasion. *Working Paper*, 2017b.
- Elliot Lipnowski and Doron Ravid. Cheap talk with transparent motives. *Econometrica*, 88(4):1631–1660, 2020.
- Elliot Lipnowski, Doron Ravid, and Denis Shishkin. Persuasion via weak institutions. *Available at SSRN 3168103*, 2019.
- George J Mailath and Larry Samuelson. *Repeated games and reputations: long-run relationships*. Oxford university press, 2006.

- Chiara Margaria and Alex Smolin. Dynamic communication with biased senders. *Games and Economic Behavior*, 110:330–339, 2018.
- Laurent Mathevet, David Pearce, and Ennio Stacchetti. Reputation and information design. *Working Paper*, 2019.
- Jeffrey Mensch. Monotone persuasion. *Games and Economic Behavior*, 2021.
- Roger B Myerson. Multistage games with communication. *Econometrica*, 52(2):323–358, 1986.
- Chris Nosko and Steven Tadelis. The limits of reputation in platform markets: An empirical analysis and field experiment. Technical report, National Bureau of Economic Research, 2015.
- Eduardo Perez-Richet. Interim Bayesian persuasion: First steps. *American Economic Review*, 104(5):469–474, 2014.
- Eduardo Perez-Richet, Vasiliki Skreta, et al. Test design under falsification. Technical report, CEPR Discussion Papers, 2021.
- Roy Radner. Repeated Principal-Agent Games with Discounting. *Econometrica*, 53(5):1173–1198, 1985.
- Luis Rayo and Ilya Segal. Optimal information disclosure. *Journal of Political Economy*, 118(5):949–987, 2010.
- Herman Rubin and Oscar Wesler. A note on convexity in euclidean n-space. In *Proceedings of the American Mathematical Society*, volume 9, pages 522–523, 1958.
- Walter Rudin. *Principles of mathematical analysis*, volume 3. McGraw-hill New York, 1964.
- Andrés Salamanca. The value of mediated communication. *Journal of Economic Theory*, 192:105191, 2021.
- Takuo Sugaya and Alexander Wolitzky. Maintaining privacy in cartels. *Journal of Political Economy*, 126(6):2569–2607, 2018.
- Steven Tadelis. Reputation and feedback systems in online platform markets. *Annual Review of Economics*, 8:321–340, 2016.

Appendix A: Definitions & Proofs of Preliminaries

In the case of complete public records, the sender’s strategy $\sigma : \mathcal{H} \times \Theta \rightarrow \Delta M$, receivers’ strategies $\rho : \underline{\mathcal{H}} \times M \rightarrow \Delta A$, and beliefs $\mu_h : \underline{\mathcal{H}} \times M \rightarrow \Delta \Theta$ are measurable functions of the respective histories.⁴¹; to emphasize it is a function, we subscript $\mu(\cdot, \cdot)$ by h . Profile $\langle \sigma, \rho \rangle$ induces a distribution \mathbb{P} over histories h_t in the usual way. To emphasize the dependence of \mathbb{P} and expectations, \mathbb{E} , on $\langle \sigma, \rho \rangle$, we may write $\mathbb{P}_{\langle \sigma, \rho \rangle}, \mathbb{E}_{\langle \sigma, \rho \rangle}$. The sender’s strategy is *public* if it can be expressed as $\underline{\sigma} : \underline{\mathcal{H}} \times \Theta \rightarrow \Delta M$. We write $\underline{\sigma}$ below where we feel the need to emphasize a strategy is public. In what follows, we denote $|\Theta| = N$.

Equilibrium

Here we define a *public* PBE (PPBE) for complete records.⁴² PBE in which σ is not public play little role in our analysis of complete records (see Lemma 2), so we leave their obvious definition to the reader. A PPBE is a triple $\langle \sigma, \rho, \mu_h \rangle$ in which σ is public and

1. Sender optimality: Given ρ , σ satisfies

$$\sigma \in \arg \max_{\underline{\sigma}} \mathbb{E}_{\langle \underline{\sigma}, \rho \rangle} \left[\sum_{\tau=1}^{\infty} \delta^{\tau-1} v(a_{\tau}) \mid h_t \right], \quad \forall h_t \in \mathcal{H}.$$

2. Receiver optimality: Given $\langle \sigma, \mu_h \rangle$, ρ satisfies

$$\rho \left(\arg \max_{a \in A} \sum_{\theta_t \in \Theta} u(\theta_t, a_t) \mu_h(\theta_t \mid \underline{h}_t, m_t) \mid \underline{h}_t, m_t \right) = 1, \quad \forall (\underline{h}_t, m_t) \in \underline{\mathcal{H}} \times M.$$

3. Consistent beliefs: For $\underline{h}_t \in \underline{\mathcal{H}}$ \mathbb{P} -a.e., $\theta' \in \Theta$ and Borel subsets $\hat{M} \subseteq M$:

$$\sum_{\theta_t \in \Theta} \int_{\hat{M}} \mu_h(\theta' \mid \underline{h}_t, m) d\sigma(m \mid \underline{h}_t, \theta_t) \mu_0(\theta_t) = \mu_0(\theta') \sigma(\hat{M} \mid \underline{h}_t, \theta').$$

Conditions 1 - 2 require players best respond at the relevant histories; 3 extends the condition that beliefs obey Bayes’ rule where possible to potentially non-discrete histories.⁴³

⁴¹ \underline{h}_t also contains a history of public randomizations y_t . For ease, we suppress it when possible.

⁴²Kuvalekar et al. (2021) give a related definition in a repeated cheap talk model without monitoring.

⁴³That is, μ_h a.e. equals the relevant Radon-Nikodym derivative.

Direct Equilibrium (DE)

In the below, it will be useful to have separate notation for the direct game. We write m^d for messages in $\Delta\Theta$. Let $h_t^d := (m_\tau^d, y_\tau, a_\tau, \omega_\tau)_{\tau=1}^{t-1}$ denote a *direct history* of the direct game, \mathcal{H}_t^d the set of time- t histories, and \mathcal{H}^d the set of all histories.⁴⁴ We identify ‘**direct strategies**’ with (measurable) functions $\lambda^d : \mathcal{H}^d \rightarrow \Delta\Delta\Theta$, $\rho^d : \mathcal{H}^d \times \Delta\Theta \rightarrow \Delta A$ and beliefs with $\mu^d : \mathcal{H}^d \times \Delta\Theta \rightarrow \Delta\Theta$. Note that $\lambda^d(h_t^d)$ refers to a probability measure on $\Delta\Theta$ induced at h_t^d , while for any Borel B , $\lambda^d(B | h_t^d)$ yields a corresponding probability; ρ^d, μ^d are treated similarly. To aid comparison with Bayesian Persuasion, $\lambda^d(h_t^d)$ is expressed as a marginal, but should be understood to derive from conditional distributions $\lambda^d(\cdot | h_t^d, \theta_t)$ (see condition 7 below, for example). Formally, DE is a triple $\langle \lambda^d, \rho^d, \mu^d \rangle$ satisfying:

4. Sender optimality: Given ρ^d , λ^d satisfies

$$\lambda^d \in \arg \max_{\tilde{\lambda}^d} \mathbb{E} \langle \tilde{\lambda}^d, \rho^d \rangle \left[\sum_{\tau=1}^{\infty} \delta^{\tau-1} v(a_\tau) | h_t^d \right], \quad \forall h_t^d \in \mathcal{H}^d.$$

5. Receiver optimality: Given $\langle \lambda^d, \mu^d \rangle$, ρ^d satisfies

$$\rho^d \left(\arg \max_{a \in A} \sum_{\theta \in \Theta} u(\theta, a) \mu^d(\theta | h_t^d, m_t^d) \mid h_t^d, m_t^d \right) = 1, \quad \forall (h_t^d, m_t^d) \in \mathcal{H}^d \times \Delta\Theta.$$

6. Bayes plausibility (BP): For all h_t^d ,

$$\int_{\Delta\Theta} m_t^d d\lambda^d(m_t^d | h_t^d) = \mu_0.$$

7. Obedience: For $h_t^d \in \mathcal{H}^d$ \mathbb{P} -a.e., $\theta' \in \Theta$, and Borel $\hat{\Delta}\Theta \subset \Delta\Theta$

$$\sum_{\theta_t \in \Theta} \int_{\hat{\Delta}\Theta} \mu^d(\theta' | h_t^d, m_t^d) d\lambda^d(m_t^d | h_t^d, \theta_t) \mu_0(\theta_t) = \mu_0(\theta') \lambda^d(\hat{\Delta}\Theta | h_t^d, \theta').$$

Moreover, $\mu^d(h_t^d, m_t^d) = m_t^d$ for any $m_t^d \in \text{supp } \lambda^d(h_t^d)$.

Conditions 4 - 5 express best response criteria. 6 further requires that the sender always induces Bayes plausible distributions over messages *in equilibrium*; To consistency, 7 adds that the receiver’s beliefs must match messages in $\text{supp } \lambda(h_t^d)$ (but not for $m_t^d \notin \text{supp } \lambda(h_t^d)$).

Lemma 1

We prove Lemma 1 in four steps: Lemma 2 justifies a focus on public PBE. Given any PPBE, Lemma 3 then identifies the induced information structures as a measurable function of the public history. As this function is not defined on direct histories, we perform a relabeling which does yield appropriate replicating strategies and beliefs for the direct game. Lemma 4 verifies that these establish the ‘only if’ direction of Lemma 1; Lemma 5 proves the converse.

⁴⁴Note the randomization y_t realizes *after* m_t ; this is not crucial for our results, but aids exposition.

Lemma 2 *For any PBE, there is a PPBE which induces the same marginal distributions over $(\theta_t, a_t, \omega_t)$, $\forall t$.*

Since ρ is public and θ_t is i.i.d., the sender's equilibrium strategy can always be modified by replacing any θ^t -dependent randomization with its conditional expectation (given \underline{h}_t), without affecting the sender's payoff. This modification is public and obviously does not change incentives. As the proof is just an application of the arguments in Theorem 5.2 Fudenberg and Levine (1994), we omit it.⁴⁵ The next Lemma, which extends a result of Kamenica and Gentzkow (2011) to our repeated setting, is a key step in connecting PBE and DE. Since μ_h is a measurable map, $\mu_h^{-1}(B)$ is a Borel subset of $\underline{\mathcal{H}} \times M$ for any Borel $B \subset \Delta\Theta$. For the next result, we make use of the conditional inverse map $\mu_h^{-1}(B | \underline{h}_t) = \{m \in M : \mu_h(\underline{h}_t, m) \in B\}$.

Lemma 3 *For any public strategy $\underline{\sigma} : \underline{\mathcal{H}} \times \Theta \rightarrow \Delta M$, there exists a measurable function $\lambda : \underline{\mathcal{H}} \rightarrow \Delta\Delta\Theta$ such that $\forall \underline{h}_t \in \underline{\mathcal{H}}$,*

$$\int_{\Delta\Theta} m_t^d d\lambda(m_t^d | \underline{h}_t) = \mu_0, \quad (4)$$

and, for any Borel subset $B \subseteq \Delta\Theta$,

$$\lambda(B | \underline{h}_t) = \sum_{\theta_t \in \Theta} \underline{\sigma}(\mu_h^{-1}(B | \underline{h}_t) | \underline{h}_t, \theta_t) \mu_0(\theta_t), \quad \underline{h}_t\text{-a.e.} \quad (5)$$

Proof: Since $\Delta\Theta$ is a Polish space, there exists (Theorem 44.3, Bauer (2011)) a probability kernel $\lambda : \underline{\mathcal{H}} \rightarrow \Delta\Delta\Theta$, a measurable function such that $\lambda(\underline{h}_t)$ satisfies

$$\int_H \lambda(B | \underline{h}_t) d\mathbb{P}(\underline{h}_t) = \mathbb{P}[B \times H] \quad (6)$$

for each Borel $B \subset \Delta\Theta$ and $H \subset \mathcal{H}_t$, \mathbb{P} -a.e. To establish (5), we show $\nu(\cdot | \underline{h}_t) := \sum_{\theta_t \in \Theta} \underline{\sigma}(\mu_h^{-1}(\cdot | \underline{h}_t) | \underline{h}_t, \theta_t) \mu_0(\theta_t)$ —the distribution of posteriors induced at \underline{h}_t by $\underline{\sigma}$ —is a valid version of λ . As $\Delta\Theta$ is Polish, we need only show $\nu(B | \underline{h}_t)$ is measurable- \underline{h}_t for each Borel B (Theorem 44.3, Bauer (2011)), and satisfies (6) almost everywhere. Fixing some Borel $B \subset \Delta\Theta$, measurability follows from an argument almost identical to that in Billingsley (1995) (pg. 232). To see that it satisfies (6), apply the law of iterated expectations to the indicator $\mathbb{1}((\underline{h}_t, m_t) \in \mu_h^{-1}(B))$, recognizing that $\mathbb{P}[\hat{M} | \underline{h}_t] = \sum_{\theta_t \in \Theta} \underline{\sigma}(\hat{M} | \underline{h}_t, \theta_t) \mu_{0,\theta_t}$ holds for any Borel $\hat{M} \subset M$. Finally, note that ν is a measure on $\Delta\Theta$ (Billingsley (1995), pg. 185), corresponding to the distribution of the random variable $m_t^d = \mu_h(\underline{h}_t, m_t)$. The change-of-variables formula applied to the left-hand side of condition 3 for $\hat{M} = M$ shows (4) also holds at each \underline{h}_t . \square

We cannot yet draw from Lemma 3 the conclusions of Lemma 1. This is because λ is a function of \underline{h}_t , which is not a direct history. Lemma 4 below constructs a ‘well-behaved’ relabeling of the histories \underline{h}_t , strategies σ, ρ and beliefs μ_h which will form an appropriate replicating DE. To make our argument, we first need to define some additional objects.

⁴⁵Measurability of the modified strategy follows trivially from the finiteness of θ^t . Details on request.

In an equilibrium $\langle \sigma, \rho, \mu_h \rangle$, messages may not only influence current beliefs; they can also coordinate continuation play. However, in a DE, messages correspond only to recommended beliefs. Nonetheless, a public randomization device (p.r.d) can perform the coordination. To this end, we encode messages with an invertible, measurable mapping $f : M \rightarrow [0, 1]$.⁴⁶ In a PPBE, each $f(m_t)$ is a random variable whose distribution derives from the equilibrium measure $\mathbb{P}_{\langle \sigma, \rho \rangle}$. By another application of Theorem 44.3, Bauer (2011), there exist kernels $\varphi : \underline{\mathcal{H}} \times \Delta\Theta \rightarrow \Delta[0, 1]$, $\bar{\varphi} : \underline{\mathcal{H}} \rightarrow \Delta[0, 1]$, which satisfy, for any Borel $\hat{Y} \subseteq [0, 1]$, $B \subseteq \Delta\Theta$

$$\mathbb{P}_{\langle \sigma, \rho \rangle} \left[(f(m_t), \theta_t) \in \hat{Y} \times B \mid \underline{h}_t \right] = \int_B \varphi \left(\hat{Y} \mid \underline{h}_t, m_t^d \right) d\lambda \left(m_t^d \mid \underline{h}_t \right) \quad (7)$$

\mathbb{P} -a.e., and $\bar{\varphi}(\underline{h}_t) = \int_{\Delta\Theta} \varphi \left(\hat{Y} \mid \underline{h}_t, m_t^d \right) d\lambda \left(m_t^d \mid \underline{h}_t \right)$. These are the conditional distributions of $f(m_t)$ given (\underline{h}_t, m_t^d) and \underline{h}_t , respectively, where $m_t^d = \mu_h(\underline{h}_t, m_t)$. With a small abuse, we set $y_t = f(m_t)$ below in our replicating DE. As will become clear below, y_t can be considered a device on which continuation play is coordinated. From the view of our replicating DE, the original PPBE is fixed and hence y_t is exogenous; indeed, in the DE below, the sender will be able to choose m_t^d but not y_t directly. As we describe it, y_t is still not a traditional p.r.d.—its distribution varies with \underline{h}_t . However, we emphasize our argument requires only that y_t has the same *distribution* as $f(m_t)$. Because of this, y_t can be considered a (history-dependent) function of a true, i.i.d. uniform p.r.d. (see for instance Billingsley (1995), Theorem 14.1), and hence we simply identify it with the latter; indeed, m_t plays no explicit role in the DE.

We also introduce a function $d : \underline{\mathcal{H}} \rightarrow \mathcal{H}^d$ that replaces each history in $\underline{\mathcal{H}}$ with an appropriate direct history, defined inductively as follows: put $d(\underline{h}_1) = \emptyset$ and for $t > 1$,

$$d(\underline{h}_t) = (d(\underline{h}_{t-1}), \mu_h(\underline{h}_{t-1}, m_{t-1}), y_{t-1}, a_{t-1}, \omega_{t-1}).$$

At each t , d updates the direct public history to include the realized time $t-1$ beliefs, actions and feedback. As f and μ_h are measurable, so too is d .⁴⁷ Notice also that d is invertible.

Not all histories in \mathcal{H}^d need live in $d(\underline{\mathcal{H}})$.⁴⁸ However, to define a replicating DE, we must specify direct strategies as a function of the *full set* \mathcal{H}^d . To deal with this, we introduce a ‘projection’ from \mathcal{H}^d into $d(\underline{\mathcal{H}})$. Define a (measurable) state variable $\psi : \mathcal{H}^d \rightarrow \mathcal{H}^d$ (on which strategies will depend) by setting $\psi(h_1^d) = h_1^d = \emptyset$ and, for $t > 1$,

$$\psi \left(h_t^d \right) = \begin{cases} (\psi(h_{t-1}^d), m_{t-1}^d, y_{t-1}, a_{t-1}, \omega_{t-1}), & \text{if } m_{t-1}^d \in \text{supp } \lambda(d^{-1}(\psi(h_{t-1}^d))); \\ (\psi(h_{t-1}^d), \mu_h(d^{-1} \circ \psi(h_{t-1}^d), f^{-1}(y_{t-1})), y_{t-1}, a_{t-1}, \omega_{t-1}), & \text{otherwise.} \end{cases}$$

ψ recursively updates histories by (i) recording the belief corresponding to the latest message, if and only if those correspond to the sender’s strategy σ at history $\underline{h}_{t-1} = d^{-1}(\psi(h_{t-1}^d))$, and otherwise (ii) replacing the sender’s chosen message m_t^d with the one corresponding (i.e., from the PPBE) to the realized draw of y_t . As Lemma 4 makes clear, the role of ψ is to specify off-path beliefs and play. By definition, $\psi(h_t^d) = h_t^d$ if and only if $m_\tau^d \in \text{supp } \lambda(d^{-1}(h_{\tau-1}^d))$ for all $\tau \leq t-1$. In other words, the set on which $\psi(h_\tau^d) = h_\tau^d$ for all $\tau \leq t$ is exactly $d(\underline{\mathcal{H}}_t)$.

⁴⁶Since M is Polish, f exists and has a measurable inverse (by the Borel isomorphism theorem).

⁴⁷By the $\pi - \lambda$ theorem, we need only see that d^{-1} maps measurable rectangles in \mathcal{H}^d into Borel sets. Letting $H_t^d \subset \mathcal{H}_t^d$ be such a rectangle, this is clear for H_1^d and extends by induction to H_t^d , $t > 1$.

⁴⁸If σ induces λ without full support at \underline{h}_t , there is $m^d \in \Delta\Theta$ such that $(d(\underline{h}_t), m^d, y_t, a_t, \omega_t) \notin d(\underline{\mathcal{H}})$.

To see that ψ is measurable, suppose $\psi(h_\tau^d)$ is measurable for all $\tau \leq t-1$ (this is clear for $t=2$).⁴⁹ Then by the usual composition rule, $\lambda(d^{-1}(\psi(h_\tau^d)))$ is a measurable function of h_τ^d . Along with a similar construction for ρ , these functions induce a measure \mathbb{P}' on \mathcal{H}_t^d .⁵⁰ As the above discussion makes clear, $d(\underline{\mathcal{H}}_t)$ is the support of this distribution and hence a Borel set, for each t . But ψ is an identity map on $d(\underline{\mathcal{H}}_t)$ and, given measurability of μ_h and f^{-1} , is a measurable function of history on $d(\underline{\mathcal{H}}_t)^c$. Hence by induction, ψ is measurable. We can now replicate a PPBE of the original game with a DE:

Lemma 4 *For any public PBE, there exists a corresponding DE which induces the same marginal distributions over $(\theta_t, a_t, \omega_t)$, $\forall t$.*

Proof: We construct a DE with the required properties. For the sender's strategy, put $\lambda^d = \lambda \circ d^{-1} \circ \psi$, for the receiver's strategy, $\rho^d = \rho(d^{-1}(\psi(h_t^d)), f^{-1}(y_t))$ and for beliefs, $\mu^d = \mu_h(d^{-1}(\psi(h_t^d)), f^{-1}(y_t))$.⁵¹ These functions correspond in an obvious way to mappings of the form $\lambda^d: \mathcal{H}^d \rightarrow \Delta\Delta\Theta$, $\rho^d: \mathcal{H}^d \times \Delta\Theta \times [0, 1] \rightarrow \Delta A$, $\mu^d: \mathcal{H}^d \times \Delta\Theta \rightarrow \Delta\Theta$. Moreover, by the usual composition rule all these maps are measurable. Endow y_t with the conditional distribution $\phi: \mathcal{H}^d \times \Delta\Theta \rightarrow \Delta[0, 1]$ where for any Borel $\hat{Y} \subseteq [0, 1]$:

$$\phi(\hat{Y} \mid h_t^d, m_t^d) = \begin{cases} \varphi(\hat{Y} \mid h_t^d, m_t^d), & \text{if } m_t^d \in \text{supp } \lambda^d(h_t^d), \\ \bar{\varphi}(\hat{Y} \mid h_t^d), & \text{otherwise.} \end{cases}$$

By pinning down the evolution of y_t we complete the description of a stochastic process for h_t^d . In particular, $\langle \lambda^d, \rho^d, \phi \rangle$ induces a distribution over \mathcal{H}^d ; call it $\mathbb{Q}_{\langle \lambda^d, \rho^d, \phi \rangle}$. With the mapping between PBE and DE, it is easy to verify that for any Borel subsets $\hat{M} \subseteq M$, $\hat{A} \subseteq A$, $\theta \in \Theta$, $\omega \in \Omega$ and t , the distributions \mathbb{Q} , \mathbb{P} satisfy

$$\mathbb{P}_{\langle \sigma, \rho \rangle}[\theta_t = \theta, m_t \in \hat{M}, a_t \in \hat{A}, \omega_t = \omega] = \mathbb{Q}_{\langle \lambda^d, \rho^d, \phi \rangle}[\theta_t = \theta, y_t \in f(\hat{M}), a_t \in \hat{A}, \omega_t = \omega].$$

For $\hat{M} = M$, this reduces to $\mathbb{P}_{\langle \sigma, \rho \rangle}[\theta_t = \theta, a_t \in \hat{A}, \omega_t = \omega] = \mathbb{Q}_{\langle \lambda^d, \rho^d, \phi \rangle}[\theta_t = \theta, a_t \in \hat{A}, \omega_t = \omega]$. In other words, \mathbb{Q} and \mathbb{P} agree on all the measurable rectangles in $\Theta \times A \times \Omega$. Hence, by Dynkin's $\pi - \lambda$ theorem, they agree on every Borel subset of $\Theta \times A \times \Omega$, establishing that $\langle \lambda^d, \rho^d, \phi \rangle$ and $\langle \sigma, \rho \rangle$ induce the same distribution over $(\theta_t, a_t, \omega_t)$, for all t .

Last, we verify $\langle \lambda^d, \rho^d, \mu^d \rangle$ augmented by the p.r.d. y_t constitutes a DE. Since $\langle \lambda^d, \rho^d, \mu^d, \phi \rangle$ and $\langle \sigma, \rho \rangle$ induce identical distributions over $(\theta_t, a_t, \omega_t)$, the sender's payoff from using λ^d is exactly the same as under $\langle \sigma, \rho \rangle$. Moreover, it is easy to see that, by construction, he has the same set of deviations (up to relabelling) in the DE as in PBE. Hence, λ^d satisfies condition 4. By Lemma 3, it is clear that $\lambda^d(h_t^d)$ satisfies (4) and hence condition 6, for each $h_t^d \in \mathcal{H}^d$. Moreover, by definition of μ^d it is clear that for any $h_t^d \in d(\underline{\mathcal{H}})$, $m_t^d \in \text{supp } \lambda^d(h_t^d)$, $\mu^d(h_t^d, m_t^d) = \mu_h(d^{-1}(\psi(h_t^d)), f^{-1}(y_t)) = \mu_h(d^{-1}(h_t^d), f^{-1}(y_t)) = \mu_h(\underline{h}_t, m_t) = m_t^d$, where

⁴⁹The justification for this inductive approach is the same as that in footnote 47.

⁵⁰Strictly, we must also define a process for y_t ; we use the one defined in the proof of Lemma 4.

⁵¹The careful reader may worry that ρ^d and μ^d depend on y_t rather than m_t^d . However—by the definition of ϕ below—message m_t^d induces only values of y_t consistent with those m_t inducing posterior m_t^d in the PPBE, so the receiver's beliefs are exactly pinned down by the sender's message.

the last equality follows from the final part of the proof of Lemma 3 . Hence, the same is true for λ^d at the appropriate direct history $h_t^d = d^{-1}(\psi(\underline{h}_t))$. Condition 7 follows by a change of variables applied to 3, using $m_t^d = \mu_h(\underline{h}_t, m_t) = \mu^d(h_t^d, m_t^d)$. Finally, it is obvious that ρ^d satisfies condition 5, given that (i) $m_t^d = \mu_h(\underline{h}_t, m_t)$ for $m_t^d \in \text{supp } \lambda^d(h_t^d)$ and ρ obeys condition 2, and (ii) for $m_t^d \notin \text{supp } \lambda^d(h_t^d)$, (μ^d, ρ^d) replicate those at some PPBE history. \square

Conversely, the following Lemma closes the equivalence between PBE and DE:

Lemma 5 *For any DE $\langle \lambda^d, \rho^d, \mu^d \rangle$, there exists an equilibrium $\langle \sigma, \rho, \mu_h \rangle$ of the game with message space M which induces the same marginal distributions over $(\theta_t, a_t, \omega_t)$, $\forall t$.*

If the sender's message space were $\Delta\Theta$, Lemma 5 would be an easy extension of KG, so we omit the details.⁵² Since M is Polish and uncountable, there is a measurable, one-to-one map between M and $\Delta\Theta$ from which relabeling the result follows. \square

Proposition 1

First, we show that it is without loss for payoffs to consider only equilibria in which there are always at most $N + 1$ messages in the support of the sender's strategy. Given a sequence $(a_\tau)_{\tau=1}^\infty$ let $V = (1 - \delta) \sum_{\tau=1}^\infty \delta^{\tau-1} v(a_\tau)$, so that $\mathcal{V}_\lambda(h_t^d) = \mathbb{E}[V | h_t^d]$ is the sender's equilibrium expected discounted payoff given h_t^d . Similarly, given sequence $(a_\tau, \theta_\tau, \omega_\tau)_{\tau=1}^\infty$ define 'receiver welfare' by $U = (1 - \delta) \sum_{\tau=1}^\infty \delta^{\tau-1} u(\theta_\tau, a_\tau)$, with corresponding equilibrium value $\mathcal{U}_\lambda(h_t^d)$.⁵³

Lemma 6 *For any DE $\langle \lambda^d, \rho^d, \mu^d \rangle$, there is another, $\langle \tilde{\lambda}^d, \tilde{\rho}^d, \tilde{\mu}^d \rangle$, in which (1.) $|\text{supp } \tilde{\lambda}^d(h_t^d)| \leq N + 1$, $\forall h_t^d \in \mathcal{H}^d$, and (2.) $\mathcal{V}_{\tilde{\lambda}^d}(h_1^d) = \mathcal{V}_{\lambda^d}(h_1^d)$, $\mathcal{U}_{\tilde{\lambda}^d}(h_1^d) = \mathcal{U}_{\lambda^d}(h_1^d)$.*

Proof: Set $\tilde{\rho}^d = \rho^d$, $\tilde{\mu}^d = \mu^d$. To find $\tilde{\lambda}^d$, we first show that for any $h \in \mathcal{H}^d$ and corresponding $\lambda^d(h)$, there exists an alternate BP information structure λ'_h with three properties: (i) $\text{supp } \lambda'_h \subset \text{supp } \lambda^d(h)$; (ii) $|\text{supp } \lambda'_h| \leq N + 1$; and (iii) $\mathcal{V}(h)$, $\mathcal{U}(h)$ are unchanged if λ^d is adjusted by replacing $\lambda^d(h)$ with λ'_h at h only. As $\lambda^d(h)$ is BP, it follows from Rubín and Wesler (1958) that $(\mu_0, \mathcal{V}(h), \mathcal{U}(h)) \in \text{co}(X_h)$, where $X_h = \{(\mu, u, v) : u = \mathbb{E}[U | h, \mu], v = \mathbb{E}[V | h, \mu], \mu \in \text{supp } \lambda^d(h)\}$. Hence, there exists a *finite* subset $\{x_h^s\}_{s=1}^k \subset X_h$ (where k may depend on h) and corresponding (probability) weights $(\lambda_h^s)_{s=1}^k$ satisfying (iii).⁵⁴ From here, the existence of λ'_h satisfying (i)-(iii) follows from Carathéodory's Theorem.⁵⁵

⁵²Following KG, it is trivial to construct (state-dependent) strategies which induce the desired information structures. The only additional step here is to verify the strategy is measurable, which follows straightforwardly from another application of Theorem 44.3, Bauer (2011).

⁵³ \mathcal{V} and \mathcal{U} also depend on ρ : we suppress this dependence as ρ is fixed in the proof of Lemma 6.

⁵⁴Recall $\text{co}(X)$ is equal to the set of all convex combinations of finitely many points in X .

⁵⁵Carathéodory's Theorem alone reduces the support to at most $N + 2$ -points; we show in a previous version (available on request) that the sender's incentive compatibility allows a further reduction by 1.

Given this observation, we now construct a sequence of auxiliary functions $\tilde{\lambda}_t^d : \mathcal{H}^d \rightarrow \Delta\Delta\Theta$, $t = 1, 2, \dots$, each measurable. Set $\tilde{\lambda}_1^d$ such that $\tilde{\lambda}_1^d(h_1^d) = \lambda'_{h_1^d}$, and $\tilde{\lambda}_1^d(h) = \lambda^d(h)$ for all $h \in \mathcal{H}^d \setminus \{h_1^d\}$. Then inductively define for all $t > 1$:

$$\tilde{\lambda}_t^d(h) = \begin{cases} \tilde{\lambda}_\tau^d(h), & \text{for } h \in \text{supp } \tilde{\mathbb{P}}_\tau, \tau < t, \\ \lambda'_h, & \text{for } h \in \text{supp } \tilde{\mathbb{P}}_t, \\ \lambda_\tau^d(h), & \text{for } h \in \text{supp } \tilde{\mathbb{P}}_\tau, \tau > t. \end{cases}$$

where $\tilde{\mathbb{P}}_t$ is the distribution over \mathcal{H}_t^d induced by $\tilde{\lambda}_{t-1}^d$. $\tilde{\lambda}_t^d(h)$ adjusts λ^d only for $\tau < t$, replacing its information structures with ones that always have at most an $N + 1$ -point support. With such a strategy, the set of time t -histories can be treated as finite.⁵⁶ Moreover, by construction, $\mathcal{V}_{\lambda^d}(h_1^d) = \mathcal{V}_{\tilde{\lambda}_t^d}(h_1^d)$ and $\mathcal{U}_{\lambda^d}(h_1^d) = \mathcal{U}_{\tilde{\lambda}_t^d}(h_1^d)$, for all t . Taking $t \rightarrow \infty$, these auxiliary functions define a function $\tilde{\lambda}^d : \mathcal{H}^d \rightarrow \Delta\Delta\Theta$ such that $\tilde{\lambda}^d(h_t^d) = \lambda'_{h_t^d}$ for all t . This direct strategy obviously satisfies properties (i) and (ii) above. Moreover, it is easy to verify that, due to discounting, $\mathcal{V}_{\lambda^d}(h_1^d) = \mathcal{V}_{\tilde{\lambda}^d}(h_1^d)$ and $\mathcal{U}_{\lambda^d}(h_1^d) = \mathcal{U}_{\tilde{\lambda}^d}(h_1^d)$ and hence, property (iii) holds. Verifying equilibrium is simple: the receiver clearly still best responds given beliefs, and from KG, λ' can be constructed to obey conditions 6 and 7 for all h . Finally, as the sender's equilibrium payoff is unchanged and he has weakly fewer deviations, condition 4 holds. \square

We now turn to the properties of $\mathcal{E}(\delta)$.⁵⁷ Considering Lemmas 1 and 6, we may take M to be finite in what follows. We show that our sequential-move stage game does not introduce any novel issues for characterizing equilibrium payoffs. In particular, the set of PPBE payoffs coincide with a set of Perfect Public Equilibrium payoffs (Mailath and Samuelson (2006), pg. 231). Define the set of actions justified by belief μ as

$$\bar{A}(\mu) := \{a \in A : a \in \arg \max_{\mathbb{E}_\mu} [u(\theta, a)]\},$$

and let $\bar{A} := \bigcup_{\Delta\Theta} \bar{A}(\mu)$ be the set of *justifiable actions*. We note that \bar{A} is a compact set. \bar{A} is closed because $\bar{A}(\mu)$ is nonempty, compact-valued, and upper hemicontinuous (by the Theorem of the Maximum). As a closed subset of the compact set A , it is also compact.

Denote the repeated game of Section 1, in which the sender's message space is M and the receiver's action space A , by $G_{(M,A)}$. We compare the PPBE of $G_{(M,A)}$ to the PPE of $G_{(M,\bar{A})}$.⁵⁸ A Perfect Public Equilibrium of $G_{(M,\bar{A})}$ is a profile of public strategies $\langle \sigma, \rho \rangle$, where $\sigma : \mathcal{H} \times \Theta \rightarrow \Delta M$ and $\rho : \mathcal{H} \times M \rightarrow \Delta \bar{A}$ are measurable and obey equilibrium conditions 1 and 2, above. Define the associated set of PPE payoffs in this adjusted game as

$$\mathcal{E}^*(\delta) := \left\{ (u, v) : \exists \text{ a PPE } \langle \sigma, \rho \rangle \text{ of } G_{(M,\bar{A})} \text{ s.t. } u = \mathbb{E}_{\langle \sigma, \rho \rangle} [U], v = \mathbb{E}_{\langle \sigma, \rho \rangle} [V] \right\}. \quad (8)$$

Of course, every PPBE $\langle \sigma, \rho, \mu_h \rangle$ in $G_{(M,A)}$ induces a PPE strategy profile $\langle \sigma, \rho \rangle$ of $G_{(M,\bar{A})}$. The next Lemma verifies the converse.

⁵⁶While we are not explicit about the output of $\tilde{\lambda}_t$ for off-path histories, we implicitly 'complete' the description of these direct strategies with ψ , exactly as in the proof of Lemma 4.

⁵⁷As μ_0 remains constant in our arguments here, we suppress the dependence of \mathcal{E} on it.

⁵⁸In fact, both are equivalent to the PPE payoffs of $G_{(M,A)}$ too. We take our approach because it allows us to describe off-path strategies in Lemma 7 particularly simply.

Lemma 7 For any PPE $\langle \sigma, \rho \rangle$ of repeated game $G_{(M, \bar{A})}$, there exists a PPBE of $G_{(M, A)}$ $\langle \sigma', \rho', \mu_h \rangle$ in which: $\sigma' = \sigma$ for all $\underline{h}_t \in \underline{\mathcal{H}}$ and $\rho' = \rho$ for all (\underline{h}_t, m) such that $m \in \text{supp } \sigma(\underline{h}_t)$.

Proof: As \bar{A} is compact and v continuous, $\arg \min_{\bar{A}} v(a)$ is nonempty; let \check{a} be one of its members, and $\check{\mu}$ a belief justifying \check{a} . Fix a PPE $\langle \sigma, \rho \rangle$ of $G_{(M, \bar{A})}$, and define ρ' as follows: for any (\underline{h}_t, m) such that $m_t \in \text{supp } \sigma(\underline{h}_t)$, set $\rho'(\underline{h}_t, m_t) = \rho(\underline{h}_t, m_t)$; otherwise, $\rho'(\{a_t = \check{a}\} | \underline{h}_t, m_t) = 1$.⁵⁹ For any (\underline{h}_t, m) such that $m \in \text{supp } \sigma(\underline{h}_t)$, set beliefs

$$\mu_h(\theta | \underline{h}_t, m_t) = \frac{\mu_0(\theta) \sigma(m_t | \underline{h}_t, \theta)}{\sum_{\theta' \in \Theta} \mu_0(\theta') \sigma(m_t | \underline{h}_t, \theta')},$$

for each $\theta \in \Theta$, and set $\mu_h(\theta | \underline{h}_t, m_t) = \check{\mu}_\theta$, for each $\theta \in \Theta$ otherwise.⁶⁰ Clearly, $\langle \sigma, \rho \rangle$ and $\langle \sigma, \rho' \rangle$ induce identical distributions on $\underline{\mathcal{H}}$. We show $\langle \sigma, \rho', \mu_h \rangle$ constitutes a PPBE of $G_{(M, A)}$. μ_h obviously satisfies consistency by construction. Consider receiver optimality: ρ' satisfies condition 2 for any $m \in \text{supp } \sigma(\underline{h}_t)$ —by definition of ρ —and trivially for any $m_t \notin \text{supp } \sigma(\underline{h}_t)$. For sender optimality, note that since $\rho' = \rho$ whenever $m_t \in \text{supp } \sigma(\underline{h}_t)$, we have

$$\mathbb{E}_{\langle \sigma, \rho \rangle} \left[\sum_{\tau=1}^{\infty} \delta^{\tau-1} v(a_t) \mid \underline{h}_t \right] = \mathbb{E}_{\langle \sigma, \rho' \rangle} \left[\sum_{\tau=1}^{\infty} \delta^{\tau-1} v(a_t) \mid \underline{h}_t \right].$$

Moreover, since at any (\underline{h}_t, m_t) ρ' is either identical to ρ or induces an action in $\arg \min_{\bar{A}} v(a)$, any alternative strategy $\tilde{\sigma}$ must yield $\mathbb{E}_{\langle \tilde{\sigma}, \rho' \rangle} [v(a_t) | \underline{h}_t] \leq \mathbb{E}_{\langle \sigma, \rho' \rangle} [v(a_t) | \underline{h}_t]$ for each \underline{h}_t (recall the receiver's action space is \bar{A} in $G_{(M, \bar{A})}$). Taking discounted sums and then expectations:

$$\mathbb{E} \left[\sum_{\tau=1}^{\infty} \delta^{\tau-1} v(a_t) \mid \underline{h}_t; \sigma, \rho' \right] \geq \mathbb{E} \left[\sum_{\tau=1}^{\infty} \delta^{\tau-1} v(a_t) \mid \underline{h}_t; \tilde{\sigma}, \rho' \right]. \quad \square$$

Lemma 7 implies $\mathcal{E}(\delta) = \mathcal{E}^*(\delta)$. We now argue that $\mathcal{E}^*(\delta)$ is amenable to the analysis of Abreu et al. (1990), and as a result, compact. To do that, we first note that it is sufficient to consider equilibria in which any receiver mixes over at most a bounded number of actions:

Lemma 8 Any $(u, v) \in \mathcal{E}^*(\delta)$ can be attained with a PPE $\langle \sigma, \rho \rangle$ s.t. $|\text{supp } \rho(\underline{h}_t)| \leq |\Theta| + 1, \forall \underline{h}_t$.

The proof is very close to that of Lemma 6—using a combination of Carathéodory's Theorem (fixing sender-type θ_t 's continuation payoffs to each message, and hence his incentives) and a similar sequential method for adjusting the receiver's strategy—so we omit the details.

Lemma 9 $\mathcal{E}(\delta)$ is compact, convex, and increasing in δ (in the set inclusion order)

Proof: By Lemma 8, consider w.l.o.g. PPE in which the range space of ρ is $(A^{|\Theta|+1} \times \mathbb{R}^{|\Theta|})^M$.⁶¹ As M is finite, A is compact metrizable (and hence sequentially compact) and we

⁵⁹ \check{a} acts as a *credible punishment* for 'off-path' deviations by the sender under strategy ρ' .

⁶⁰As M is finite, it is easy to check ρ', μ are measurable (and indeed, $\tilde{\rho}, \rho''$).

⁶¹i.e., for each $m \in M$, ρ specifies a distribution with $|\Theta| + 1$ -point support on A .

allow for public randomization, standard results (see for instance Abreu et al. (1990), and chapter 7.3, Mailath and Samuelson (2006)) imply that $\mathcal{E}^*(\delta)$ is a compact, convex set, and is increasing in δ in the set inclusion order. While their methods apply to the payoffs of long-run players only, they can be extended to \mathcal{E}^* by adding a fictitious long-run player who takes no action but has the payoffs U . Lemma 7 and Lemma 1 establish the result for $\mathcal{E}(\delta)$. \square

Appendix B: Proofs of Main Results

Theorem 1

We first show $V_{de}(\mu_0) = V_{kg}(\mu_0)$ if and only if $V_{kg}(\mu_0) = \tilde{v}(\mu_0)$. The ‘if’ direction is trivial, as $\tilde{v}(\mu_0) \leq V_{de}(\mu_0) \leq V_{kg}(\mu_0)$ always holds (see Lemma 12). For the converse, we use the following well-understood property of V_{kg} :

Lemma 10 *Fix $\mu_0 \in \Delta\Theta$. Let λ^{kg} solve (KG) and denote $\text{supp } \lambda^{kg} = \{\mu^{*,k}\}_{k=1}^K$. For any $\mu = \sum_{k=1}^K \alpha^k \mu^{*,k}$, where $\alpha_k \geq 0$, $k = 1, \dots, K$, and $\sum_{k=1}^K \alpha_k = 1$,*

$$V_{kg}(\mu) = \sum_{k=1}^K \alpha^k v(\mu^{*,k}). \quad (9)$$

Proof: By definition, $V_{kg}(\mu) \geq \sum_{k=1}^K \alpha^k v(\mu^{*,k})$. Towards a contradiction, suppose the inequality is strict for some $\mu \in \text{co } \{\mu^{*,k}\}_{k=1}^K$. As $\lambda_k^{kg} := \lambda^{kg}(\mu^{*,k}) > 0$ for all k , there is a $\mu' \in \text{co } \{\mu^{*,k}\}_{k=1}^K$ and $\gamma \in (0, 1)$ such that $\mu_0 = \gamma\mu + (1-\gamma)\mu'$ (see footnote 62). Then

$$\gamma V_{kg}(\mu) + (1-\gamma)V_{kg}(\mu') > \gamma \sum_{k=1}^K \alpha^k V_{kg}(\mu^{*,k}) + (1-\gamma) \sum_{k=1}^K \beta^k V_{kg}(\mu^{*,k}) = \sum_{k=1}^K \lambda^{kg}(\mu^{*,k}) \cdot V_{kg}(\mu^{*,k}),$$

where $\{\beta^k\}_{k=1}^K$ denotes convex weights associated with μ' , chosen to satisfy $\lambda_k^{kg} = \gamma\alpha^k + (1-\gamma)\beta^k$ for all k .⁶² The RHS of this chain is at least $V_{kg}(\mu_0)$, contradicting concavity of V_{kg} . \square

We now prove that $V_{de}(\mu_0) = V_{kg}(\mu_0)$ implies $V_{kg}(\mu_0) = \tilde{v}(\mu_0)$. If $V_{de}(\mu_0) = V_{kg}(\mu_0)$, then from (CP) there must exist some $\lambda^{kg} \in \Lambda_{kg}(\mu_0)$ such that $c(\lambda^{kg}) = 0$. This implies that either (i) λ^{kg} is partitional, or (ii) for every $\theta \in \Theta$ and $\mu \in B_\theta := \text{supp } \lambda^{kg}(\cdot | \theta)$, $v(\mu) = \underline{v}_\theta(\lambda^{kg})$. Obviously we need only consider case (ii). Fix some $\theta \in \Theta$ for which B_θ is not a singleton. By Lemma 10, $V_{kg}(\mu) = \underline{v}_\theta(\lambda^{kg})$ for all $\mu \in \text{co } B_\theta$. Furthermore, as $v(\mu)$ is constant on B_θ , the information structures in (9) can be supported as one-shot cheap talk equilibria and so $V_{ct}(\mu) = V_{kg}(\mu)$ for all $\mu \in \text{co } B_\theta$. Under Assumption 1, this implies $v(\mu) = \underline{v}_\theta$ for all $\mu \in \text{co } B_\theta$.

As payoffs are constant on B_θ , each θ can be identified with a unique payoff \underline{v}_θ . Index these payoffs by i (a small but useful abuse of notation: i need not take consecutive integer values here), partition Θ according to these payoffs—that is, let $\theta, \theta' \in P_i$ if and only if $\underline{v}_\theta = \underline{v}_{\theta'} = i$ —and notice that for every $\theta \in P_i$ and $\mu \in B_\theta$ we have $\text{supp } \mu \subset P_i$ (since a μ with

⁶²Choose $\beta^k := \lambda_k^{kg} - \varepsilon(\alpha^k - \lambda_k^{kg})$, for $\varepsilon = \frac{\gamma}{1-\gamma} > 0$ small. Note $\mu' = \sum_{k=1}^K \beta^k \mu^{*,k}$ and $\sum_{k=1}^K \beta^k = 1$.

$v(\mu) = i$ is induced with positive probability only if $\theta \in P_i$). As a result, it is easy to see that the conditional expectation $\mathbb{E}[\mu \mid \mu \in \bigcup_{P_i} B_\theta]$ satisfies $\mathbb{E}[\mu \mid \mu \in \bigcup_{P_i} B_\theta] = \mathbb{E}[\mu \mid \theta \in P_i] = \mu_i$ (recall μ_i is the posterior associated with learning $\theta \in P_i$). Consider the information structure λ' which adjusts λ^{kg} by pooling posteriors $\mu \in \bigcup_{P_i} B_\theta$ into the single posterior μ_i . By the above reasoning, λ' is obviously partitional. Moreover, since $v(\mu) = i$ for all $\mu \in \bigcup_{P_i} B_\theta$, the previous paragraph shows $v(\mu_i) = i$ too and so we have $v(\lambda') = v(\lambda^{kg})$. Hence, $\tilde{v}(\mu_0) = V_{kg}(\mu_0)$.

We now show that $V_{de}(\mu_0) = V_{kg}(\mu_0)$ holds for all priors if and only if \tilde{v} is concave. Of course, if $V_{de}(\mu) = V_{kg}(\mu)$ holds for all μ then the first part of the theorem implies \tilde{v} inherits the concavity of V_{kg} . To prove the other direction, we make use of two simple observations:

Lemma 11 *Consider two real functions f, g on $\Delta\Theta$. If $f \geq g$ everywhere, then $\text{cav } f \geq \text{cav } g$.*

Lemma 11 is a well-known property of concave envelopes. Hence we omit the proof.

Lemma 12 *$v \leq \tilde{v} \leq V_{de} \leq V_{kg}$ everywhere on $\Delta\Theta$.*

Proof: The first inequality follows because ‘no information’ is a partitional information structure, the second from (CP) and $c(\lambda^P(\mu_0)) = 0$ for partitional information structures, and the third by direct comparison of (CP) and (KG). \square

We can now prove the final part of the theorem. Suppose that \tilde{v} is concave. Then $\tilde{v} = \text{cav } \tilde{v}$. Lemmas 11 and 12 imply $\text{cav } \tilde{v} = V_{kg}$, so $\tilde{v} = V_{kg}$. By Lemma 12, $V_{de} = V_{kg}$ everywhere. \square

Proposition 3

Proof: Before proving the result, we provide a useful lemma. Letting $N = |\Theta|$, $\Delta\Theta$ can obviously be identified with a subset of \mathbb{R}^{N-1} for which $\mu_\theta \geq 0$ and $\sum_{\Theta} \mu_\theta = 1$. Let \mathcal{L} be the Lebesgue measure on $\Delta\Theta$, and for a subset $P \subset \Theta$ define the vector μ_0^P by

$$\mu_{0,\theta}^P = \begin{cases} \frac{\mu_{0,\theta}}{\sum_{\theta \in P} \mu_{0,\theta}} & , \text{ if } \theta \in P \\ 0 & , \text{ otherwise.} \end{cases}$$

Lemma 13 *Suppose action set A is finite. For any partition $\mathcal{P} = \{P_i\}_{i=1}^k$, μ_0 satisfies \mathcal{L} -a.e.: (i) $\mu_0 \in \text{int } \Delta\Theta$, and (ii) $\exists \varepsilon_{\mu_0} > 0$ such that $v(\mu) = v(\mu_0^{P_i})$ for all $|\mu - \mu_0^{P_i}| < \varepsilon_{\mu_0}$, $i = 1, \dots, k$.*

Proof: The sets $E(\theta) := \{\mu_0 \in \Delta\Theta : \mu_{0,\theta} = 0\}$ are each contained in $N - 2$ dimensional hyperplanes in \mathbb{R}^{N-1} , so that $\mathcal{L}(E(\theta)) = 0$ and hence $\mathcal{L}(E) = 0$, where $E = \bigcup_{\Theta} E(\theta)$. For any singleton subset $P \subset \Theta$, conditionally strict preferences trivially imply (ii). To extend this observation to all subsets $P \subset \Theta$, take any pair $a, a' \in A$ and define the set of priors $I(a, a', P) := \{\mu_0 \in E^c : du(a, a')^T \mu_0^P = 0\}$ for which a receiver would be indifferent between a and a' after learning $\theta \in P$, where $du(a, a') = (u(\theta, a) - u(\theta, a'))_{\theta \in \Theta}$. Let \mathbf{P} be an $N \times N$ matrix whose columns are indicators for members of P , and 0 otherwise. For $\mu_0 \in E^c$,

$du(a, a')^T \mu_0^P = 0 \iff (\mathbf{P}^T du(a, a'))^T \mu_0 = 0$, and by conditional strictness $\mathbf{P}^T du(a, a') \neq 0$. The latter therefore imposes a linear restriction on μ_0 , so $I(a, a', P)$ is contained in an $N - 2$ dimensional hyperplane in \mathbb{R}^{N-1} and $\mathcal{L}(I(a, a', P)) = 0$. As there are only finitely many combinations (a, a', P) , this immediately extends to $\mathcal{L}(I) = 0$, where $I = \bigcup_{A \times A \times 2^\Theta} I(a, a', P)$. Hence, $\mathcal{L}(I^c \cap E^c) = 1$. Noting $E^c = \text{int } \Delta\Theta$, the receiver's best response to any partitional information structure is strict on I^c and payoffs are continuous completes the argument. \square

We can now prove the Proposition. By Lemma 13, we can restrict attention to $\mu_0 \in E^c \cap I^c$. Obviously, if $\mu_0 \in \arg \max v$, then we are done. Suppose $\mu_0 \notin \arg \max v$ and let $\mathcal{P} = \{P_1, \dots, P_k\}$ be an optimal partition at μ_0 . Recall $\mu_i = \mu_0^{P_i}$ and $\lambda_i = \lambda^{\mathcal{P}}(\mu_i)$. For $\mu', \mu'' \in \Delta\Theta$, define $\mu'_i, \mu''_i, \lambda'_i, \lambda''_i$ in the analogous way. Define $J = \{i \in \{1, \dots, k\} : \mu_i \in \arg \max v\}$, and for an arbitrary subset $B \subset \{1, \dots, k\}$ write $\lambda_B = \sum_{i \in B} \lambda_i$. We show there exist μ', μ'' and $\alpha \in (0, 1)$ such that (i) $\alpha\mu' + (1 - \alpha)\mu'' = \mu_0$, (ii) $v(\mu'_i) = v(\mu_i)$ for $i \in J^c$ and (iii) $v(\mu'') = \max v$, $v(\mu'') > v(\mu_i)$ for some i with $\lambda''_i > 0$. With these properties, the chain

$$\begin{aligned} \tilde{v}(\mu_0) = \sum_{i=1}^k \lambda_i v(\mu_i) &= \alpha \sum_i \lambda'_i v(\mu_i) + (1 - \alpha) \sum_i \lambda''_i v(\mu_i) \\ &< \alpha v^{\mathcal{P}}(\mu') + (1 - \alpha) v(\mu'') \\ &\leq \alpha \tilde{v}(\mu') + (1 - \alpha) \tilde{v}(\mu'') \end{aligned} \quad (10)$$

is valid, and the result follows (Theorem 1). To do this, suppose first that J is nonempty. As $\mu_0 \in E^c$, $\lambda_J > 0$. Under Assumption 1, $\arg \max v$ is convex.⁶³ Hence, $\mu_0 \notin \arg \max v$ implies J^c is nonempty and $\lambda_{J^c} > 0$. Choose $\mu' = \frac{\sum_{J^c} \lambda_i \mu_i}{\lambda_{J^c}}$ and let $\tilde{\mu} = \frac{\sum_J \lambda_i \mu_i}{\lambda_J}$. We construct μ'' from μ' and $\tilde{\mu}$. Since $\mu \in I^c$, $\exists \bar{w} \in (0, 1)$ such that $w\tilde{\mu} + (1 - w)\mu' \in \arg \max v$ for $w \in (\bar{w}, 1)$. Take $\tilde{w} \in (\max\{\bar{w}, \lambda_J\}, 1)$ and set $\mu'' = \tilde{w}\tilde{\mu} + (1 - \tilde{w})\mu'$. Now, for $\alpha = 1 - \frac{\lambda_J}{\tilde{w}} \in (0, 1)$, (i) is easily verified. For (ii), we argue that $\mu'_i = \mu_i$ for $i \in J^c$. If $\theta \notin \bigcup_J P_j$ then clearly the probability μ'_θ of state θ under belief μ' is 0 and so $\mu'_{i,\theta} = 0 = \mu_{i,\theta}$ for any $i \in J^c$. For any other θ , we have $\mu'_{i,\theta} = \frac{\Pr_{\mu'}[i, \theta]}{\Pr_{\mu'}[i]} = \frac{(\lambda_i/\lambda_{J^c})\mu_{i,\theta}}{\lambda_i/\lambda_{J^c}} = \mu_{i,\theta}$. For (iii), recall that $\mu'_\theta = 0$ for all $\theta \in \bigcup_J P_j$, so for such θ (i) reduces to $(1 - \alpha)\mu''_\theta = \mu_\theta$. Therefore $(1 - \alpha)\lambda''_J = (1 - \alpha) \sum_J \sum_{P_i} \mu''_\theta = \sum_J \sum_{P_i} \mu_\theta = \lambda_J$. But, by definition of α , $\lambda_J = (1 - \alpha)\tilde{w}$. Hence, $\lambda''_J = \tilde{w} < 1$, so $\lambda''_{J^c} > 0$, which establishes (iii). The proof for the empty- J case is similar, and so omitted. \square

Proposition 4

(If) By definition, $\tilde{v} \geq v_{\mathcal{P}}$, and by (i), $v_{\mathcal{P}} \geq v$. Lemmas 11 and 12 then imply $\text{cav } v_{\mathcal{P}} = V_{kg}$. Hence if we show $v_{\mathcal{P}}$ is concave, we will be done (Theorem 1). Consider points μ, μ' and $\mu'' = \alpha\mu + (1 - \alpha)\mu'$. Let $\lambda, \lambda', \lambda''$ be the information structures corresponding to \mathcal{P} under the

⁶³Otherwise, there would be $\mu', \mu'' \in \arg \max v$ and some $\alpha \in (0, 1)$ such that $\mu = \alpha\mu' + (1 - \alpha)\mu'' \notin \arg \max v$ – an obvious violation of Assumption 1.

respective priors, and write $\gamma_j = \frac{\alpha\lambda_j}{\alpha\lambda_j + (1-\alpha)\lambda'_j}$. Then

$$\begin{aligned}
v_{\mathcal{P}}(\mu''_0) &= \sum_j \lambda''_j v(\mu''_j) \\
&= \sum_j \left(\alpha\lambda_j + (1-\alpha)\lambda'_j \right) v(\mu''_j) \\
&\geq \sum_j \left(\alpha\lambda_j + (1-\alpha)\lambda'_j \right) \left(\gamma_j v(\mu_j) + (1-\gamma_j)v(\mu'_j) \right) \\
&= \alpha v_{\mathcal{P}}(\mu) + (1-\alpha)v_{\mathcal{P}}(\mu'),
\end{aligned}$$

where line 2 uses $\lambda''_j = \alpha\lambda_j + (1-\alpha)\lambda'_j$, and line 3 uses $\mu''_j = \gamma_j\mu_j + (1-\gamma_j)\mu'_j$ and (ii). \square

(*Only if*) Suppose (i) is violated. Then there exists some μ such that $v(\mu) > v_{\mathcal{P}}(\mu)$ —contradicting optimality of \mathcal{P} . Similarly if (ii) is violated then there exists some j and $\mu \in \Delta P_j$ such that $V_{kg}(\mu) > v(\mu) = v_{\mathcal{P}}(\mu)$ since by definition v and $v_{\mathcal{P}}$ coincide on ΔP_j . \square

Proposition 5

We first define two relevant (pure) communication strategies for the stage game. Let σ^T be the *truthful* strategy, where $\sigma^T(\text{‘like new’} \mid \theta) = 1$ if $\theta = h$ and 0 otherwise, and let σ^b be ‘babbling’: $\sigma^b(\text{‘like new’} \mid \theta) = 1$ for $\theta = h, l$. We construct an equilibrium of the incomplete record game where the sender uses either σ^T or σ^b as a function of the public record, r_t . For brevity, we use strategies which from $t = 3$ onward depend on the public record indirectly via a public state variable, $s_t \in \{s^T, s^b\}$, which evolves as follows: at $t = 3$, if $\omega_1 \neq b$ then $s_t = s^T$. If $\omega_1 = b$, $\omega_2 \neq b$ then $s_3 = s^T$ with probability $1 - q \geq 0$; and if $\omega_1 = \omega_2 = b$ $s_3 = s^T$ with probability $1 - q - z \geq 0$ for parameters $q, z \geq 0$. At $t > 3$, if $s_{t-1} = s^b$ then $s_t = s^b$. However, if $s_{t-1} = s^T$ then $s_t = s^T$ with certainty if $\omega_t \neq b$ and with probability $1 - z$ if $\omega_t = b$.

Consider the following strategy profile. The sender plays σ^T if $t = 1$, if $t = 2$ and $\omega_1 = b$, and if $s_t = s^T$, $t \geq 3$. Otherwise, he plays σ^b . If $t \in \{1, 2\}$ or $s_t = s^T$, then receiver buys if and only if $m_t = \text{‘like new’}$. Otherwise she never buys, regardless of m_t . Finally set $q = (1 - \delta) \frac{1 - \delta(p + \mu_0 - 2p\mu_0)}{\delta^2 \mu_0 (2p - 1)}$, $z = \frac{1 - \delta}{\delta \mu_0 (2p - 1)}$ and $l = -1$. Note $q, z \geq 0$, since $p + \mu_0 - 2\mu_0 p = p(1 - \mu_0) + (1 - p)\mu_0$ and $p > \frac{1}{2}$. Moreover for $p > \frac{1}{2}$, $\mu_0 \in (0, 1)$, q, z are continuous in δ , with $q = z = 0$ at $\delta = 1$. Thus there is a threshold $\underline{\delta}(p, \mu_0) < 1$ s.t. $q + z \leq 1$ for $\delta \geq \underline{\delta}(p, \mu_0)$.

We first verify equilibrium. Trivially each receiver at $t \neq 2$ is playing a best response. It is also easy to show the sender’s strategy is a best response for $\delta \geq \underline{\delta}(p, \mu_0)$.⁶⁴ At $t = 2$, the receiver’s best response is indeed to buy iff $m_2 = \text{‘like new’}$, so long as:

$$\Pr[\theta = h \mid m_2 = \text{‘like new’}] = \frac{\mu_0}{1 - (1 - \mu_0) \Pr[\omega_1 = b]} = \frac{\mu_0}{1 - (1 - \mu_0)\mu_0(1 - p)} \geq \frac{1}{2}. \quad (11)$$

For any $p > \frac{1}{2}$, the LHS is continuous in μ_0 and (11) holds strictly at $\mu_0 = \frac{1}{2}$. Thus for μ_0 close enough to $\frac{1}{2}$ the second receiver best responds too.

⁶⁴Available on request.

Let $V^T = \frac{1}{1-\delta}\mu_0(1 - \frac{1-p}{p})$. Applying Fudenberg and Levine (1994), it is easy to see V^T is the sender's maximal payoff in any equilibrium with a complete public record. We now show players can be better off on average. The sender's payoff in this equilibrium can be written

$$V' = V^T + \delta \left(1 - \mu_0 + \mu_0 \frac{1-p}{p} \right) > V^T. \quad (12)$$

The discounted sum of receivers' utility can be written

$$U' = U^T + \delta \left(\mu_0 \frac{1-p}{p} - (\Pr[\omega_1 \neq b] - \Pr[\omega_1 = b]) (1 - \mu_0) |l| \right), \quad (13)$$

where $\Pr[\omega_1 = b] = \mu_0(1-p)$, $U^T = \frac{1}{1-\delta}\mu_0(1 - \frac{1-p}{p})$. On rearrangement, $U' > U^T$ iff

$$\frac{1-p}{p} > (1 - 2\mu_0(1-p)) \frac{1-\mu_0}{\mu_0}. \quad (14)$$

Clearly both LHS and RHS are continuous in p, μ_0 , for $\mu_0 > 0$. At $\mu_0 = \frac{1}{2}$, (14) is satisfied strictly for $\frac{1}{2} \leq p < \frac{\sqrt{5}-1}{2}$. Thus clearly there exist μ_0, p close enough to $\frac{1}{2}$ to satisfy both (11) and (14). Indeed both hold for $p = \frac{6}{11}$, $\mu_0 = \frac{5}{11}$. Finally, for any such choice this is an equilibrium so long as $1 > \delta \geq \underline{\delta}(p, \mu_0)$. For $p = \frac{6}{11}$, $\mu_0 = \frac{5}{11}$, $\delta \geq \frac{123}{126}$ is sufficient. \square

Theorem 2

Throughout, we denote SBS phase lengths by $\Gamma_G = \Gamma$ and $\Gamma_B = \beta \times \Gamma$ respectively. We also take the sender's message space M as finite, with $|M| = 2N + 2$; smaller spaces can be incorporated by assigning 0 mass to some messages (see footnote 69), while larger spaces will not be necessary.⁶⁵ As a result, we may identify $\Lambda(\mu_0)$ with an appropriate subset of $\mathbb{R}^{M(N+1)}$, endowed with the Euclidean topology.⁶⁶ For an arbitrary $\lambda \in \Lambda(\mu_0)$ we index members μ_m , $m = 1, \dots, M$, of $\text{supp } \lambda$ by the messages that induce them, and write $\lambda_m = \lambda(\mu_m)$ (recall partitional information structures are the special case where μ_i is induced iff $P_i \in \mathcal{P}$).

Fix some $(u', v') \in \mathcal{CS}$. We show that for any $\varepsilon > 0$ there is a threshold $\delta_\varepsilon < 1$ such that when $\delta \geq \delta_\varepsilon$ a simple badge system (SBS) and corresponding equilibrium exist with payoffs (u, v) satisfying $\|(u, v) - (u', v')\| \leq \varepsilon$. Taking $\varepsilon \rightarrow 0$ yields the result. First, we introduce some useful preliminaries.

Preliminary: a continuity consequence of Assumption 3

Following the argument of Proposition 1, Kamenica and Gentzkow (2011), there exists λ^{**} such that $(u(\lambda^{**}), v(\lambda^{**})) = (u', v')$ and $|\text{supp } \lambda^{**}| \leq N+2$. We identify another helpful information structure λ^* that induces the same distribution as λ^{**} , but may additionally specify some

⁶⁵In fact, the results hold for $|M| \leq N + 2$, at the cost of additional steps to prove Lemma 14.

⁶⁶Any such $\lambda \in \Lambda(\mu_0)$ maps naturally to an $(N+1) \times M$ matrix with columns $\begin{bmatrix} \mu \\ \lambda(\mu) \end{bmatrix}_m$ for $m \in M$.

‘redundant’ (0 probability) posteriors. As the next result shows, λ^* is useful because it provides a degree of freedom to choose approximating information structures λ^n whose payoffs converge smoothly to $(u(\lambda^*), v(\lambda^*)) = (u', v')$, without violating Bayes plausibility:

Lemma 14 *Suppose Assumption 3 holds and $\mu_0 \in \text{int } \Delta\Theta$. For any $\lambda^{**} \in \Lambda(\mu_0)$ such that $|\text{supp } \lambda^{**}| \leq N + 2$, there exists $\lambda^* \in \Lambda(\mu_0)$ with $\text{supp } \lambda^* = \text{supp } \lambda^{**}$, a sequence $\lambda^n \rightarrow \lambda^*$ and open balls $B^n(\lambda^n)$ such that (i) $v(\lambda)$ is continuous on $\text{cl}(B_n(\lambda^n))$, (ii) $\lambda^* \in \text{cl}(\bigcup_n B_n(\lambda^n))$ and (iii) $v(\lambda^n) \rightarrow v(\lambda^*)$.*

Proof: Let $M = \{1, \dots, 2N + 2\}$, and index $\text{supp } \lambda^{**}$ by $m \in \{1, \dots, N + 2\}$. We construct a sequence λ^n which converges to some λ^* with the desired properties. To do so, fix a sequence $\varepsilon_n \downarrow 0$, and for each n choose a set of ‘approximating posteriors’ $\mu_m^n \in B_{\varepsilon_n}(\mu_m^{**}) \cap X_{\mu_m^{**}}$, where $B_\varepsilon(\mu)$ is the open ε -ball around μ (such posteriors exist by Assumption 3). Of course, if we assign the probabilities $\{\lambda_m^{**}\}_{m=1}^{N+2}$ to $\{\mu_m^n\}_{m=1}^{N+2}$, Bayes plausibility (BP) may be violated; write $z^n = \|\sum_{m=1}^{N+2} \lambda_m^{**}(\mu_m^{**} - \mu_m^n)\|$ for the ‘size’ of this violation, and notice $z^n \leq \varepsilon_n, \forall n$. To recover BP, we introduce an additional posterior ν^n to each approximating set as follows: choose some $\bar{x} > 0$ such that the closed \bar{x} -ball $\text{cl}(B_{\bar{x}}(\mu_0)) \subset \text{int } \Delta\Theta$ (\bar{x} exists as $\mu_0 \in \text{int } \Delta\Theta$), and let $\gamma^n = \frac{\bar{x}}{\bar{x} + z^n}$. We choose $\nu^n = \mu_0 + \frac{\gamma^n}{1 - \gamma^n} \sum_{m=1}^{N+2} \lambda_m^{**}(\mu_m^{**} - \mu_m^n)$. It is easily verified that ν^n has the following properties: (i) $\nu^n \in \text{cl}(B_{\bar{x}}(\mu_0)), \forall n$, (ii) $\mu_0 = \gamma^n \sum_{m=1}^{N+2} \lambda_m^{**} \mu_m^n + (1 - \gamma^n)\nu^n$, and (iii) $\gamma^n \rightarrow 1$.

If each ν^n were known to be a continuity point of v , the result would follow for the sequence of obvious information structures on $\text{supp } \lambda^{**} \cup \{\nu^n\}$. As we cannot guarantee this, we augment these information structures by replacing ν^n with a BP lottery over continuity points of v as follows: index Θ by $j = 1, \dots, N$ and for $m = N + 2 + j$ let e_m be the basis vector for state j . Then trivially, $\nu^n \in \text{int co } \{e_m\}_{m=N+2+j}^{2N+2}$. As the e_m are linearly independent, there is a real sequence $w_n \downarrow 0$ such that $\mu_m \in B_{w_n}(e_m), \forall m \in \{N + 3, \dots, 2N + 2\}$, implies $\nu^n \in \text{int co } \{\mu_m\}_{m=N+3}^{2N+2}$ (see Rudin (1964), Theorem 9.8(b) and p.211). By Assumption 3, for each $m = N + 3, \dots, 2N + 2$ we can therefore choose $\mu_m^n \in B_{w_n}(e_m) \cap X_{e_j}$, and convex weights α_m^n such that $\nu^n = \sum_{m=N+3}^{2N+2} \alpha_m^n \mu_m^n$. Let λ^n be the obvious information structure supported on $\{\mu_i^n\}_{i=1}^{N+2} \cup \{\tilde{\mu}_j^n\}_{j=1}^N$.

To conclude, note that λ^n has a subsequence which converges to some λ^* with the desired properties (Bolzano-Weierstrass theorem). The required λ^* sets $\mu_m^* = \mu_m^{**}$, for $m \leq N + 2$ and $\mu_m^* = e_m$ otherwise. For each n along this subsequence, $\mu_m^n \in X_{\mu_m^{**}}$ or $\mu_m^n \in X_{e_m}, \forall m \in M$, and so there is a real sequence $\varepsilon'_n \downarrow 0$ such that $B_{\varepsilon'_n}(\mu_m^n) \subset X_{\mu_m^*}, \forall m$. Choose $B^n(\lambda^n)$ to have radius ε'_n . Consider any $\lambda \in B^n(\lambda^n)$, with support $\{\mu_m\}_{m=1}^{2N+2}$. Then, since $\|\lambda - \lambda^n\|^2 = \sum_{m=1}^{2N+2} |\lambda_m - \lambda_m^n|^2 + \sum_{m=1}^{2N+2} \|\mu_m - \mu_m^n\|^2$, $\lambda \in B^n(\lambda^n)$ implies $\mu_m \in B_{\varepsilon'_n}(\mu_m^n), \forall m \in M$. Hence, each $v(\mu_m)$ is continuous on $B_n(\lambda^n)$; therefore so too is $\sum_{m=1}^{2N+2} \lambda_m v(\mu_m)$, from which the result follows. \square

Preliminary: λ -Tight Standards

Let $h'_l(h_t) = (\theta_\tau, m_\tau, a_\tau, \omega_\tau)_{\tau=t-l}^{t-1}$ be the sub-component of h_t formed from the most recent $l \leq t$ periods ($h'_0 = \emptyset$), and $\|\cdot\|_n$ the Euclidean metric on \mathbb{R}^n (we drop the n subscript when

clear from context). We consider a class of SBS which use the following standards. For some ‘target’ $\lambda \in \Lambda(\mu_0)$, evaluation length Γ and a ‘tolerance’ $\chi > 0$, set

$$\mathcal{S} = \left\{ h'_\Gamma : \|\ell_\Gamma(m, \omega) - q(m, \omega)\| \leq \begin{cases} \chi, & \text{if } m \in M = \text{supp}(\lambda), \\ 0, & \text{otherwise.} \end{cases}, \forall \omega \in \Omega \right\}. \quad (15)$$

where $q(m, \omega) := \sum_{\Theta} p(\omega | \theta) \lambda(m | \theta) \mu_{0, \theta}$ is the joint probability of the pair m and ω implied by information structure λ . Recall from (3) that ℓ_Γ is applied to h'_Γ . We label the messages in M by the posteriors in λ (which recall is a $2N + 2$ -point distribution). Despite this labeling, we emphasize that \mathcal{S} depends *only* on the target λ , and *not* on strategies σ, ρ or beliefs μ_m .

Preliminary: Perfect Bayesian equilibrium with an SBS

We consider PBE in ‘ r -recursive’ strategies. For receivers, an r -recursive (pure) strategy $\rho : \{\mathcal{G}, \mathcal{B}\} \times M \rightarrow A$ and beliefs $\mu : \{\mathcal{G}, \mathcal{B}\} \times M \rightarrow \Delta\Theta$ depend only on the current badge status. For the sender, an r -recursive (behavior) strategy defines functions

$$\sigma^r : \bigcup_{l=1}^{\Gamma_r} \mathcal{H}_l \times \Theta \rightarrow \Delta M, \quad r = \mathcal{G}, \mathcal{B}. \quad (16)$$

In a r -recursive strategy, the sender’s behavior $\sigma^r(h'_l(h_t), \theta_t)$ in the $(l + 1)^{th}$ period t of a phase r depends on at most (i) the current badge status, r , and (ii) his private history *within the current phase*. We write $\sigma^r(m | h'_l, \theta)$ for the probability of sending m given such a history. Of course, the mixed strategy profile $\langle \sigma, \rho \rangle$ induces a distribution over h'_Γ . We note that this distribution is linear in the mixed strategy σ in the following sense (see for example Fudenberg and Tirole (1991), sections 1.1.1 and 3.4.2): for two strategies σ and σ' , the mixture $\tilde{\sigma} = \gamma\sigma + (1 - \gamma)\sigma'$ induces the distribution $\text{Pr}_{\tilde{\sigma}}[h'_\Gamma] = \gamma \text{Pr}_\sigma[h'_\Gamma] + (1 - \gamma) \text{Pr}_{\sigma'}[h'_\Gamma]$ over h'_Γ . Finally, we write μ^r for receivers’ belief system as a function of the record.

If receivers play r -recursive strategies, the sender has a best response which is r -recursive (and vice versa). A PBE in r -recursive strategies is a collection $\langle \sigma^r, \rho^r, \mu^r \rangle_{r \in \{\mathcal{G}, \mathcal{B}\}}$, where (i) given ρ, σ^r maximizes $\mathbb{E}_{\langle \tilde{\sigma}, \rho \rangle} [\sum \delta^{t-1} v(a_t) | h_t, r^t]$; (ii) given σ and μ^r, ρ^r maximizes $\sum_{\Theta} \mu^r(\theta | m) u(\theta, a)$ for each (r, m) , and finally (iii) given σ and ρ , beliefs satisfy Bayes’ rule, where possible: for each (r, m) such that $\mathbb{P}_{\langle \sigma, \rho \rangle}[m | r] > 0$,

$$\mu^r(\theta | m) = \frac{\sum_{l=1}^{\Gamma} \mathbb{P}_{\langle \sigma, \rho \rangle}[\theta, m, l]}{\sum_{\theta' \in \Theta} \sum_{l=1}^{\Gamma} \mathbb{P}_{\langle \sigma, \rho \rangle}[\theta', m, l]}, \quad (17)$$

where $\mathbb{P}_{\langle \sigma, \rho \rangle}[\theta, m, l]$ is the equilibrium probability that the receiver arrives in the l^{th} period of the current evaluation, the state at that time is θ , and the sender sends message m .⁶⁷

We prove our result in r -recursive trigger strategies: if $r_t = \mathcal{B}$, the sender’s worst stage

⁶⁷Note Bayes’ rule requires that in (17) we use the distribution over (h'_l, m) induced by the sender’s strategy and the processes for θ, ω . Hence, ‘correct beliefs’ about h'_l are implied.

Nash equilibrium is played, whose stage payoffs we normalize to $(0, 0)$. For profile $\langle \sigma^{\mathcal{G}}, \rho^{\mathcal{G}}, \mu^{\mathcal{G}} \rangle$, we can write the sender's payoff $V_{\mathcal{G}}$ as follows:

$$V_{\mathcal{G}}(\sigma^{\mathcal{G}}, \mu^{\mathcal{G}}) = \sum_{l=0}^{\Gamma-1} \delta^l \mathbb{E}_{\sigma^{\mathcal{G}}} [v(\mu^{\mathcal{G}}(m_l))] + \delta^{\Gamma} \left(\mathbb{P}_{\sigma} [r_{\Gamma} = \mathcal{G}] + \delta^{\beta\Gamma} \mathbb{P}_{\sigma} [r_{\Gamma} = \mathcal{B}] \right) V_{\mathcal{G}}(\sigma^{\mathcal{G}}, \mu^{\mathcal{G}}). \quad (18)$$

where $\Pr_{\sigma} [r_{\Gamma} = \mathcal{B}] = \sum_{\mathcal{S}} \Pr_{\sigma} [h'_{\Gamma}] = \sum_{\mathcal{S}} \prod_{l=1}^{\Gamma-1} \mu_{0,\theta} \sigma(m_l | \theta_l) p(\omega_l | \theta_l)$, θ_l, m_l and ω_l being the obvious projections of h'_l . Note that $\Pr_{\sigma} [r_{\Gamma} = \mathcal{B}]$ is continuous in the mixed strategy σ . For ease, we present the case where ties are broken in the sender's favor, allowing us to drop explicit mention of $\rho^{\mathcal{G}}$ and write stage payoffs $v(\mu^{\mathcal{G}})$; the same argument holds more generally.

Proof.

Unlike repeated games with two long-run players, we cannot use future threats to incentivize receivers. Furthermore, immediate appeal to fixed point arguments is not helpful, since a 'bad' equilibrium (one-shot cheap talk) always exists. Finally, we cannot leverage contraction mapping arguments (our result shows (18) cannot be a contraction). Instead, we show players' best responses can be trapped in a mutually closed subspace which yields the two properties described in the theorem as $\delta \rightarrow 1$; we apply existence arguments to this restricted domain.

Fix some $\lambda' \in \underline{\Delta}(\mu_0)$ such that $\|v(\lambda') - v(\lambda^*)\| < \frac{\varepsilon}{2}$, $\|u(\lambda') - u(\lambda^*)\| < \frac{\varepsilon}{2}$ and $m \in X_{\mu_m^*}$ for $m \in M$ (such a λ' exists by Lemma 14, and continuity of u); in what follows, λ' will be used as the 'target' information structure in \mathcal{S} . Further, by appropriate definition of M we may take $\lambda(m) > 0$ for all $m \in M$ (see proof of Lemma 14). To ease the flow, we assume $a(\mu) \neq a'$ for all $\mu \in \text{supp } \lambda'$, where a' is unmonitored; at the end, we describe how the proof extends.

Verifying equilibrium in \mathcal{B} -phases is trivial and thus omitted. We focus on \mathcal{G} -phases from now, and so drop ' r ' superscripts where the meaning is clear. Given the sender's (r -recursive) strategy σ , it will be useful to define the time-average strategy during a \mathcal{G} -phase as

$$\bar{\sigma}(m, \theta) := \frac{\sum_{l=1}^{\Gamma} \mathbb{E}_{\langle \sigma, \rho \rangle} [\sigma(m | h'_l, \theta) | r = \mathcal{G}]}{\Gamma}, \quad (19)$$

where the expectation is taken over h'_l , for $l = 1, \dots, \Gamma$.⁶⁸ We look for the existence of an equilibrium in which strategies and beliefs are bounded 'close to' those implied by λ' . To that end, consider for any $\varepsilon_{\mathcal{S}}, \varepsilon_R > 0$ the sets of strategies and beliefs satisfying the bounds:

$$\|\bar{\sigma}(m, \theta) - \lambda'(m | \theta)\| \leq \varepsilon_{\mathcal{S}}, \quad \forall m \in M, \theta \in \Theta; \quad \|\mu^{\mathcal{G}}(m) - m\|_N \leq \varepsilon_R, \quad \forall m \in M, \quad (20)$$

where recall $N = |\Theta|$. We refer to σ, μ satisfying (20) as $\varepsilon_{\mathcal{S}}$ -bounded, ε_R -bounded, respectively; the set of such pairs is clearly compact and convex. Throughout, we consider values of ε_R such that $\varepsilon_R \leq \varepsilon_{\lambda'}$, where $\varepsilon_{\lambda'} > 0$ is small enough that (i) $v(\lambda)$ is continuous on a closed $\varepsilon_{\lambda'}$ -ball around λ' (and so $v(\mu)$ is continuous on a $\varepsilon_{\lambda'}$ -ball around $m \in X_{\mu_m^*}$ for all $m \in \text{supp } \lambda'$), and (ii) the sender's payoff from choosing $\sigma(\cdot | \theta) = \lambda'(\cdot | \theta)$ for all θ ,

$$\min_{\{(\mu(m))_M : \|\mu(m) - m\| \leq \varepsilon_R\}} \mathbb{E}_{\lambda'} [v(\mu(m))] > 0, \quad (21)$$

⁶⁸As θ_t is independent of h_t , there is no need to condition this expectation on θ_t .

exceeds 0 when beliefs are $\epsilon_{\lambda'}$ -bounded (such $\epsilon_{\lambda'} > 0$ exists by Lemma 14 and $v(\lambda') > 0$). Furthermore, we also focus throughout on strategies which are $\epsilon_{\lambda'}$ -bounded. As a result, we additionally require $\epsilon_{\lambda'}$ be small enough that if $\lambda'(m | \theta) > 0$ for some (m, θ) then $\bar{\sigma}(m, \theta) > 0$ for all $\epsilon_{\lambda'}$ -bounded strategies. Notice that the $\epsilon_{\lambda'}$ required to bound strategies *and* beliefs in these ways depends only on λ' .

When strategies and beliefs are $\epsilon_{\lambda'}$ -bounded, (17) and (18) become amenable to fixed-point analysis. In particular, express (18) as

$$V_{\mathcal{G}}(\sigma, \mu) = \frac{s(\sigma, \mu)}{1 - \delta^{\Gamma} - \kappa \Pr_{\sigma}[r_{\Gamma} = \mathcal{B}]} \quad (22)$$

where $s(\sigma, \mu) = \sum_{l=0}^{\Gamma-1} \delta^l \mathbb{E}_{\sigma}[v(\mu(m_l))]$ and $\kappa = \delta^{\Gamma}(1 - \delta^{\beta\Gamma}) > 0$. Then it is easy to see that $\epsilon_{\lambda'}$ -bounded beliefs imply $s(\sigma, \mu)$ is continuous in (σ, μ) . Recalling that $\Pr_{\sigma}[r_{\Gamma} = \mathcal{B}]$ is continuous in σ , $V_{\mathcal{G}}$ is clearly also continuous (the denominator above being bounded away from 0). Moreover, $V_{\mathcal{G}}$ is quasiconcave in σ . To see this, consider mixed strategies σ, σ' and $\tilde{\sigma} = \gamma\sigma + (1 - \gamma)\sigma'$ and note that (after a little algebra) $V_{\mathcal{G}}(\tilde{\sigma}, \mu)$ can be written $V_{\mathcal{G}}(\tilde{\sigma}, \mu) = wV_{\mathcal{G}}(\sigma, \mu) + (1 - w)V_{\mathcal{G}}(\sigma', \mu)$, where $w = \gamma \frac{1 - \delta^{\Gamma} - \kappa \Pr_{\sigma}[r_{\Gamma} = \mathcal{B}]}{1 - \delta^{\Gamma} - \kappa \Pr_{\tilde{\sigma}}[r_{\Gamma} = \mathcal{B}]} \in (0, 1)$. Hence, $V_{\mathcal{G}}(\tilde{\sigma}, \mu) \geq \min\{V_{\mathcal{G}}(\sigma, \mu), V_{\mathcal{G}}(\sigma', \mu)\}$, as required. Finally, when σ is $\epsilon_{\lambda'}$ -bounded, $\mu^{\mathcal{G}}$ is continuous in σ (since the denominator in 17 is bounded away from 0).

Below, we argue that for any $0 < \epsilon_R \leq \epsilon_{\lambda'}$, there exists a $\epsilon'_S > 0$ (dependent on ϵ_R) such that if σ is ϵ_S -bounded for $\epsilon_S \leq \epsilon'_S$, then beliefs are ϵ_R -bounded (step 1). We then show that for any $\epsilon_S > 0$, there exist SBS parameters Γ, χ and a $\delta_{\epsilon_S} < 1$ such that if beliefs are $\epsilon_{\lambda'}$ -bounded then the sender's best response is ϵ_S -bounded when $\delta \geq \delta_{\epsilon_S}$ (steps 2 and 3). Notice that ϵ_S can therefore be chosen independently of ϵ_R . Hence, strategies and beliefs can be mutually bounded by such ϵ_R and $\epsilon_S \leq \epsilon'_S$ for some choice of SBS parameters Γ, χ and $\delta \geq \delta_{\epsilon_S}$. Given the above properties of $V_{\mathcal{G}}, \mu^{\mathcal{G}}$, an equilibrium with ϵ_S -, ϵ_R -bounded strategies and beliefs exists (Debreu (1952), Glicksberg (1952), Fan (1952)).⁶⁹ The consequences of the theorem for payoffs and information follow by taking $\epsilon_R, \epsilon_S \rightarrow 0$.

1. Verify bounds on receivers' best responses

Fix $0 < \epsilon_R \leq \epsilon_{\lambda'}$. Existence of the required ϵ'_S follows from equation (17): for each θ , the numerator can be written $\mu_{0,\theta}\bar{\sigma}(m, \theta)$, which approaches $\mu_{0,\theta}\lambda'(m | \theta)$ as $\epsilon_S \rightarrow 0$. Since $\sum_{\Theta} \mu_{0,\theta}\lambda'(m | \theta) > 0 \forall m \in \text{supp}\{\lambda'\}$, (17) converges continuously to m —that is, there exists a $\epsilon'_S > 0$ such that if σ is ϵ_S -bounded, $\epsilon_S \leq \epsilon'_S$, then μ is ϵ_R -bounded.

2. Bounding sender's incentive to deviate from standards

Fix some $z \in (0, 1)$. In this step, we establish that the existence of SBS parameters Γ_z, β_z and

⁶⁹Strictly, this establishes Nash existence. But any such Nash equilibrium has a corresponding PBE with the same payoffs. Augment messages in $\text{supp } \lambda'$ with the obvious beliefs in a \mathcal{G} -phase. For messages not in $\text{supp } \lambda'$, assign a belief corresponding to some message in $\text{supp } \lambda'$. Given the standards, such messages cause the sender to fail an evaluation with probability 1. Because receivers cannot observe the sender's previous actions, at any history including some $m \notin \text{supp } \lambda'$, the sender's strategy in PBE must involve choosing the message that induces his favored action for the rest of the phase. As the sender could replicate this play using on-path messages for a weakly lower probability of failing, such deviations are clearly sub-optimal.

χ_z such that, if μ is $\epsilon_{\chi'}$ -bounded, then a patient enough sender is willing and able to adopt a strategy in which he secures a \mathcal{G} -rating with probability exceeding $1 - z$. This will be useful in the next step, where we translate this finding into an ϵ_S -bound on the sender's best response:

Lemma 15 *Suppose μ is $\epsilon_{\chi'}$ -bounded. Then, for any $z \in (0, 1)$, there exists SBS parameters $\Gamma_z, \beta_z, \chi_z > 0$ and a $\delta_z < 1$ such that for $\delta \geq \delta_z$ the sender's optimal strategy must satisfy*

$$\Pr_{\sigma} [h'_{\Gamma} \in \mathcal{S}_z \mid r = \mathcal{G}] \geq 1 - z, \quad (23)$$

where \mathcal{S}_z uses target λ' , $\Gamma = \Gamma_z$ and $\chi = \chi_z$ with $\Gamma_z \rightarrow \infty$, $\chi_z \rightarrow 0$ as $z \rightarrow 0$. Moreover, for these SBS strategies exist which satisfy (23) and guarantee $V_G(\sigma, \mu) > 0$.

Lemma 15 is a straightforward extension of Lemma 6.1, Radner (1985), to our communication setting and hence we omit its proof.⁷⁰ $\epsilon_{\chi'}$ -bounded beliefs ensure that, if the sender chooses $\sigma(m \mid h'_l, \theta) = \lambda'(m \mid \theta)$ for all $h'_l \in \bigcup_{l=1}^{\Gamma_z} \mathcal{H}_l$, then $V_G > 0$ (recall (21)). This is enough to ensure that a patient enough sender can be given incentives to pass with a high enough probability for the appropriately chosen SBS.

3. Bounding the sender's average strategy in a \mathcal{G} -phase

We now establish that a patient sender's best response can be appropriately ϵ_S -bounded in some SBS with two lemmas, which allow us to establish the existence of a $z > 0$ such that meeting SBS standards \mathcal{S}_z with probability $1 - z$ requires an ϵ_S -bounded strategy. Applying Lemma 15, $\delta_{\epsilon_S} = \delta_z < 1$ is the necessary patience threshold corresponding to SBS \mathcal{S}_z .

Lemma 16 *Fix $\epsilon_S > 0$. For any Γ , there exists a $\eta > 0$ such that if σ obeys*

$$\|\mathbb{E}_{\sigma} [\ell_{\Gamma}(m, \omega)] - q(m, \omega)\| \leq \eta, \quad \text{for all } m \in M, \omega \in \Omega, \quad (24)$$

then it is ϵ_S -bounded.

Proof: We show by induction that $\bar{\sigma}$ and $\mathbb{E}_{\sigma}[\ell]$ are related by a continuous bijection, which is independent of Γ . Let $\ell(m, \theta, \omega)$ be the joint frequency of triple (m, θ, ω) in a \mathcal{G} -phase. We argue that, for any strategy σ and any (m, θ, ω) ,

$$\mathbb{E}_{\sigma} [\ell(m, \theta, \omega) \mid \mathcal{G}] = p(\omega \mid \theta) \mu_{0, \theta} \bar{\sigma}(m, \theta) \quad (25)$$

and hence $\mathbb{E}_{\sigma} [\ell_{\Gamma}(m, \omega) \mid \mathcal{G}] = \sum_{\Theta} p(\omega \mid \theta) \mu_{0, \theta} \bar{\sigma}(m, \theta)$. To make the inductive argument clear, we denote explicitly the dependence of the relevant likelihood on Γ by $\ell_{\Gamma}(m, \theta, \omega)$. Clearly, (25) holds for $\Gamma = 1$. To make the inductive step, assume (25) holds for all $\Gamma \leq T$. We show it also holds for $\Gamma = T + 1$. For any likelihood function, we can write

$$\ell_{T+1}(m, \theta, \omega) = \frac{T}{T+1} \ell_T(m, \theta, \omega) + \frac{1}{T} \ell_1(m, \theta, \omega). \quad (26)$$

⁷⁰Radner (1985) analyses a repeated, binary outcome principal-agent game, in which all players are long-lived. As a result, his setting has no analogue of our steps 1 and 3.

Applying (26) to $\ell(m, \theta, \omega)$ and using the law of iterated expectations, we can write:

$$\mathbb{E}_\sigma [\ell_{T+1}(m, \theta, \omega)] = \frac{T}{T+1} \mathbb{E}_\sigma [\ell_T(m, \theta, \omega)] + \frac{1}{T+1} \mathbb{E}_\sigma [\mathbb{E}_\sigma [\ell_1(m, \theta, \omega) \mid h'_T]].$$

By the inductive hypothesis, both terms on the right-hand side can be rewritten using (25). But clearly, a similar expression to (26) can be written down for $\bar{\sigma}(m, \theta)$. Applying that condition yields the claim. Thus, for each ω , $\bar{\sigma}(m) = P^{-1} \mathbb{E}[\ell(m, \omega)]$, where P is a matrix of probabilities $p(\theta \mid \omega) \propto p(\omega \mid \theta) \mu_{0, \theta}$ —invertible under Assumption 2—and $\bar{\sigma}(m)$, $\mathbb{E}[\ell(m, \cdot)]$ are vectors. Clearly, $\bar{\sigma}(m)$ is continuous in $\mathbb{E}[\ell(m, \cdot)]$, from which existence of the requisite η follows (Theorem 9.7(a), Rudin (1964)). \square

The next Lemma relates η to the existence of the required z .

Lemma 17 *Fix $\eta > 0$. There exists a $z > 0$, with corresponding SBS standards \mathcal{S}_z given by Lemma 15, such that a strategy σ obeys (24) only if it satisfies $\Pr_\sigma[r = \mathcal{G}] \geq 1 - z$.*

Proof: Fix a SBS \mathcal{S}_z and consider any σ satisfying $\mathbb{E}_\sigma[\ell(m, \omega)] \leq q(m, \omega) - \eta$ for some (m, ω) (the other case is symmetric). By the definition of \mathcal{S}_z we have

$$\Pr_\sigma \left[h'_{\Gamma(z)} \in \mathcal{S}_z \right] \leq \Pr_\sigma [\ell(m, \omega) \geq q(m, \omega) - \chi(z)].$$

But by Markov's inequality, the right-hand expression is bounded by

$$\Pr_\sigma [\ell(m, \omega) \geq q(m, \omega) - \chi(z)] \leq \frac{\mathbb{E}_\sigma[\ell(m, \omega)]}{q(m, \omega) - \chi(z)}.$$

As $z \rightarrow 0$ this bound is no more than $1 - \frac{\eta}{q(m, \omega)}$. Clearly there exists z small enough that $1 - \frac{\eta}{q(m, \omega)} < 1 - z$ for all σ such that (24) fails. \square

Having established steps 1-3, we conclude that for target λ' and any $\varepsilon_R \leq \epsilon'_\lambda$ and $\varepsilon_S \leq \epsilon'_S$ there exist standards \mathcal{S}_z , β_z and threshold $\delta_{\varepsilon_S} < 1$ such that there is an equilibrium with ε_S -bounded $\sigma^{\mathcal{G}}$ and ε_R -bounded $\mu^{\mathcal{G}}$. Taking $\varepsilon_R, \varepsilon_S \rightarrow 0$, it should be clear that part 1 of the theorem follows directly from (20). The second part follows because for small enough $\varepsilon_R, \varepsilon_S$, payoffs may be bounded within $\frac{\varepsilon}{2}$ of $(u(\lambda'), v(\lambda'))$. This follows from by continuity of $u(\lambda)$, and applying continuity of $v(\lambda)$ for $\varepsilon_R, \varepsilon_S \leq \epsilon_{\lambda'}$ and $z \rightarrow 0$ to (22).

To extend the result for the unmonitored a' , standards \mathcal{S} should not impose any restriction on messages that may induce a' under λ' . Doing so preserves the continuity of $V_{\mathcal{G}}$ in its arguments. Moreover, since there is only a single unmonitored a' , the joint probabilities with which it is induced are pinned down by the restrictions imposed on the other messages. \square

Online Appendix A: Support Material for Section 4

Here we develop the details behind the claims made in section 4.

Section 4.1

To be concrete, we develop arguments in the context of example 1 with $l = -1$, $\mu_0 = \frac{1}{3}$, $p = 1$, and $\delta \rightarrow 1$. The threshold belief is $\underline{\mu} = 0.5$ and is induced by lying half the time when quality is low: $\lambda(0.5|\theta = l) = 0.5$. Under the assumption of random arrival, an average payoff of $V_{kg}(\mu_0) = \frac{2}{3}$ can be attained easily in equilibrium using a simple badge system with $\Gamma = 2$, $\beta = \infty$, and standards requiring no negative feedback in the second period of an evaluation. In this equilibrium the seller always lies in the first period of an evaluation and is honest in the second. Hence, it could not be an equilibrium if customers knew t — the first customer would never buy knowing the seller will lie whenever $\theta = l$. Nonetheless, when t is observable the platform could replicate the effect of random arrivals with a system that randomizes evaluation dates. We construct such a system below:

Proposition 7 *Consider example 1 with $\mu_0 < \underline{\mu}$. Expected payoff $V_{kg}(\mu_0)$ is attainable using a badge system with stochastic evaluation dates when t is public.*

Proof. Consider the following badge system with random evaluation dates. At $t = 1$, the seller is given a rating $r_1 = \mathcal{G}$. At the outset, the system determines with equal probability whether the seller’s \mathcal{G} rating will be evaluated on odd, or on even, days – and this outcome is told *privately* to the seller. If an evaluation occurs in period t , the seller retains a \mathcal{G} rating until the next evaluation if he avoids negative feedback from customer t . Otherwise, his rating is switched to \mathcal{B} forever thereafter. At any t , the incoming customer observes t and the current value of r_t only.

We argue that the following strategy profile is an equilibrium for δ large enough. Moreover the seller’s discounted average payoff converges to $\frac{2}{3}$ as $\delta \rightarrow 1$. In any evaluation, the seller always recommends ‘buy’ in the first period, irrespective of current θ_t , and recommends ‘buy’ in the second period if and only if the current $\theta_t = h$. Otherwise (if $r_t = \mathcal{B}$), the seller always recommends ‘buy’ again. Each receiver adopts the following strategy: if $r_t = \mathcal{G}$, the receiver obeys the seller’s recommendations, while if $r_t = \mathcal{B}$ she does not buy. Checking equilibrium in this case is trivial. The only deviation that really needs checking is at those histories beginning in the second period of an evaluation. By the one shot deviation property, we need only check that the seller does not wish to deviate from honesty in the second period of an evaluation. This is the case if:

$$\mu_0 + \delta V^* \geq 1 + \delta \mu_0 V^*,$$

where $V^* = \frac{1+\delta\mu_0}{1-\delta^2}$ is his continuation payoff in the event of retaining a \mathcal{G} rating. On rearrangement, this becomes

$$\delta \frac{1 + \delta \mu_0}{1 - \delta^2} \geq 1.$$

As $\delta \rightarrow 1$ the RHS of this expression becomes infinite. Hence, for δ large enough, we have verified the seller is best responding. Finally, applying l’Hôpital’s rule verifies that the discounted

average payoff $(1 - \delta)V^* \rightarrow \frac{2}{3}$.

□

What if customers must be informed of the dates on which evaluations take place? Even here it turns out the platform can still replicate the required uncertainty, so long as it adopts more complex standards that depend on *full sequences* of outcomes rather than simple averages:

Proposition 8 *Consider example 1 with $\mu_0 < \underline{\mu}$. Expected payoff $V_{kg}(\mu_0)$ is attainable using a badge system with complex standards and public deterministic evaluation dates.*

Proof. Consider the following standard. Let evaluation phases consist of T periods, where T is even. At the end of any evaluation, the seller is allowed to avoid a *permanent* \mathcal{B} rating if his outcomes show either: (i) he lied only on even days, or (ii) he lied only on odd days. If he chooses the latter, then at the end of the evaluation he faces a probability q of a permanent \mathcal{B} rating, which is designed to ensure he is indifferent between the two strategies. The required q is

$$q = \frac{1 - \mu_0}{\delta + \mu_0} \frac{(1 - \delta^T)(1 - \delta)}{\delta^T}.$$

Notice that q is decreasing in δ , with $q \rightarrow 0$ as $\delta \rightarrow 1$. Hence, there exists a threshold δ' such that q is indeed a well-defined probability, so long as $\delta \geq \delta'$.

We first argue that the strategy profile described in section 4 is an equilibrium, so long as $\delta \geq \max\{\delta', \delta''\}$, where $\delta'' < 1$ is the minimal value of x such that⁷¹

$$1 - x^T \leq x^T \cdot \frac{\mu_0 + x}{1 + x}. \quad (27)$$

Verifying that customers best respond is trivial and hence omitted. Focus instead on the seller and let his continuation value from adopting the strategy proposed in the main text be V^* . If he lies only on even days during an evaluation, he gets

$$\mu_0 + \delta + \delta^2 \mu_0 + \dots + \delta^{T-2} \mu_0 + \delta^{T-1} + \delta^T V^*.$$

Similarly, if he lies on odd days he gets

$$1 + \delta \mu_0 + \delta^2 + \dots + \delta^{T-2} + \delta^{T-1} \mu_0 + \delta^T (1 - q) V^*.$$

In order to be an equilibrium, the seller must be indifferent between these strategies. Using the maximum principle, it is easy to verify that $V^* = \frac{\mu_0 + \delta}{(1 + \delta)(1 - \delta)}$ and that indifference indeed holds for our choice of q . Moreover, any other deviation cannot be profitable for $\delta \geq \delta''$. In any evaluation, the seller's optimal deviation from his equilibrium strategy is trivially to lie in all periods. This is not profitable if

$$\frac{1 - \delta^T}{1 - \delta} \leq \delta^T V^*$$

⁷¹This minimum is well-defined: the left (right) side of the inequality is continuous and decreasing (increasing) in δ . Finally, the inequality is violated at $x = 0$ and holds strictly at $x = 1$.

which on rearrangement is just (27), at $x = \delta$. Since $\delta \geq \delta''$ this is indeed satisfied. Finally, direct calculation for $\mu_0 = \frac{1}{3}$ verifies that $\lim_{\delta \rightarrow 1} (1 - \delta)V^* = \frac{2}{3}$. \square

Section 4.2

Proof of Proposition 6

For the first claim, consider the following *blind sender* review systems. At each t , the incoming customer observes the complete history of feedback \underline{h}_t . Badges are again awarded to the seller based on evaluations of standards of length T , but now by the customers rather than a third party.⁷² The important novelty here is that the seller no longer observes the individual feedback of past customers, but only the badges. Hence his history at $t > 1$ is $(r_\tau, \theta_\tau, m_\tau, a_\tau)_{\tau=1}^{t-1}$ and at $t = 1$ it is $r_1 = \mathcal{G}$. But this maps into the framework of Abreu et al. (1991). Hence we can apply their Proposition 6 to conclude that the (pure) strategy of truth telling can be supported with an expected punishment converging to 0 as the sender gets patient (and T gets large). This can be done by adopting standards in which the seller retains a \mathcal{G} rating if and only if *every* product sold receives bad feedback.

For the final part of the Proposition, we focus on public PBE of *any* blind sender system (i.e. across all possible standards). As argued in the main text this is the relevant concept for understanding the ‘reusable punishment’ insight. But the logic of Fudenberg et al. (1990) can again be applied to this problem to show that the sender’s payoff in any PPBE is bounded by Proposition 2. In example 1, this bound corresponds to the truth telling payoff. \square

Section 4.4

We briefly illustrate the wider applicability of the results of section 3.2 with an example:

Example 3 *Replace the stage game of example 1 with the following moral hazard game. The customer chooses whether to buy, $a \in \{B, N\}$. The seller has no private information but chooses an effort level, $e \in \{h, l\}$. If the product is bought, the customer sees a noisy signal of effort $\omega_t \in \{h, l\}$, where $p(\omega_t = e|e) > 0.5$. Otherwise $\omega_t = \emptyset$. Payoffs are as follows:*

		Customer	
		B	N
Seller	h	1, 1	0, 0
	l	3, -1	0, 0

⁷²For simplicity we assume here that if a sender ever receives a \mathcal{B} rating, this is permanent. We allow for a PRD, so punishments can be made a probabilistic function of outcomes if necessary

The sets of individually rational feasible stage payoffs and payoffs attainable in equilibrium are the same as in example 1. From Theorem 2 we obtain the following.

Corollary 1 *Consider example 3. Any individually rational, feasible payoff profile is attainable as an equilibrium with some SBS as $\delta \rightarrow 1$.*

The only important difference between example 3 and 1 is what a customer learns about \underline{h}_t and t before she acts. In 3 the seller's effort teaches her nothing, whereas in 1 his message provides some information. For instance, if a receiver observes a 'used' message then she can infer she is not at a history at which the sender strictly gains from mis-selling. Since example 2 does not suffer this inference problem, the argument underpinning Theorem 2 applies *a fortiori* to classic moral hazard problems too. Indeed, it would be a relatively simple extension of Theorem 2 to show this holds in a wide class of games with a long-run player facing a sequence of short-run players.