

Persuasion for the Long Run*

James Best [†] Daniel Quigley[‡]

15th November 2017

Abstract

We examine a persuasion game where concerns about future credibility are the sole source of current credibility. A long-run sender plays a cheap talk game with a sequence of short-run receivers. Even with perfect monitoring, long-run incentives do not perfectly substitute for ex-ante commitment to reporting strategies at the stage game. We then show that different methods of augmenting or garbling history can better harness long-run incentives and expand the Pareto frontier. In particular, a ‘review aggregator’ can implement average payoffs and information structures arbitrarily close to those available under ex-ante commitment. Finally, we examine applications to e-commerce and finance.

Keywords: Bayesian Persuasion; Cheap Talk; Mechanism Design; Repeated Games.

JEL Codes: D02, D82, D83.

*We thank Vince Crawford, Johannes Hörner, Ian Jewitt, Emir Kamenica, Laurent Mathevet, Margaret Meyer, John Hardman Moore, Alessandro Pavan, Ina Taneva, Adrien Vigier, Peyton Young and Bill Zame for their feedback on the Best and Quigley [2016] version of this paper.

[†]Department of Economics and Nuffield College, University of Oxford (Email: james.best@economics.ox.ac.uk)

[‡]Department of Economics and Nuffield College, University of Oxford (Email: daniel.quigley@economics.ox.ac.uk)

The cost of being caught lying today is being ignored tomorrow. What are the limits of communication and persuasion when the desire to persuade people in the future is the only source of credibility today?

The power of these long-run incentives depends on what people observe about the historical accuracy of a person's past claims. One might naturally suppose that a complete public history will be the most efficient source of credibility. Yet, online rating systems like e-Bay, Airbnb, and Tripadvisor provide aggregate statistics in the form of star ratings. Similarly, the financial industry requires brokerages to produce aggregate statistics of their past advice. Is something lost by using such aggregates? Or is it possible that garbling history can generate more efficient communication and persuasion?

To answer these questions we develop a long-run persuasion game where a long-run sender ('he') plays a sequence of pure cheap talk games with short-run receivers (each 'she'). Each period a payoff relevant state is independently drawn. The sender observes this state and chooses a message to send to the receiver, who then chooses an action based on her beliefs. The sender has preferences only over the receiver's action. Receivers observe all prior messages and a signal about each prior state. We say monitoring is perfect when the signal is noiseless, and imperfect when the signal is noisy. Importantly, we focus on short-run receivers so the sender can only influence actions via beliefs about stage payoffs i.e. via persuasion.

Our first main result is that, even under perfect monitoring, long-run incentives do not generally substitute for the exogenous commitment introduced by Kamenica and Gentzkow [2011] (KG hereafter). We derive upper bounds on the sender's payoff in any equilibrium and moreover provide strategies that implement these bounds. We show that a patient sender can achieve the same average payoffs with long-run persuasion if and only if the optimal information structure under commitment is *partitional* i.e. a deterministic mapping from the state space to the message space. To see this, suppose the sender strictly prefers receivers to take one of two actions. A partitional information structure, such as truth-telling, can be enforced under perfect monitoring with an off-path punishment of never being believed in the future after his first lie. However, a non-partitional information structure requires that the sender mix between the 'truth' and a 'lie' in some state. Unlike persuasion with commitment such mixing requires indifference between the two messages. This indifference is achieved by punishing the sender for sending the message that induces the preferred action. Hence, there is costly on path punishment and the sender gets lower expected payoffs than in KG.

We then show the sender can do better if we use a novel mechanism that we call a 'coin and cup' to augment the history with a temporary informational asymmetry. The 'coin' is a payoff irrelevant random variable with a known distribution. Each period the coin realizes a value which is privately observed by the sender and hidden from the receiver (covered by the 'cup'). The sender then gives advice and the receiver acts. The cup

is then lifted, publicly revealing the value of the coin. Any non-partitional information structure can now be converted into a partitional one by deterministically mapping both the state and the coin realizations into messages. Under perfect monitoring there is no on path punishment and we achieve the commitment payoffs - providing a potential micro-foundation for the commitment assumption in KG.

Our second main result is that, when monitoring is imperfect, a ‘review aggregator’ can implement equilibria arbitrarily close to the payoffs and information structures feasible under commitment – dominating the complete history benchmark discussed above. Review aggregator achieves this by garbling histories. It provides a binary aggregate rating, good or bad, of the sender’s past behavior. Critically, the receivers see nothing but the rating. A ‘good phase’ lasts many periods at the end of which there is a statistical test of the sender’s reports and states. If the test is passed then there is another good phase. If the test is failed the rating switches to bad, which induces a babbling equilibrium for a fixed finite number of periods. Like Radner [1985] we use very long ‘good phases’ to statistically monitor the sender almost perfectly. However, unlike Radner [1985], we find a way of doing this with short-run players. We then characterize the sender’s optimal strategy and show that it induces an *average* information structure which ensures he nearly always passes the test. Moreover due to the garbling of history, this average information structure can be enforced by short-run receivers with almost no on-path punishment.

One might think such a review strategy can be implemented without a review aggregator - it cannot. When a *short-run* receiver observes the full history she only believes the sender’s message if the sender receives an expected punishment at least equal to the additional stage payoff of his best deviation at that history. As monitoring is imperfect this results in significant on-path punishment at every period even if the sender behaves.¹ This result follows Fudenberg and Levine [1994]: when players are short-run, payoff sets under imperfect monitoring are generally a strict subset of those under perfect monitoring.

The review aggregator tackles this problem by allowing the receivers to monitor the sender’s behavior while still being ignorant of the complete history. Our key insight is that this can be used to make receivers uncertain over their position within a review phase to the extent that they can only best respond to the average information structure. Consequently, the sender’s advice is still obeyed while simultaneously allowing him a very small amount of deterministic violations. This very small amount of flexibility over the incentive compatibility constraints of both sender and receivers allows us to expand

¹This is true for partitional and non-partitional information structures. Hence, it also holds when we have a coin and a cup. However, we will see later that the coin and cup still expands the Pareto frontier under imperfect monitoring - reducing on path punishment due to both mixing and imperfect monitoring - but it still does not recover the full commitment payoffs.

the equilibrium payoff set, reclaiming the full commitment payoffs.² Moreover, the class of aggregation rules we study allows for a tractable characterization of both sender and receiver strategies in equilibrium.

These results speak to the value of aggregate statistics in online markets. According to data from the US Bureau of the Census, sales on e-Commerce websites - such as eBay, Amazon, Alibaba, and so on - have risen from 0.5% in 1999 to almost 9% of retail sales in 2016. An important function of these trading platforms is that buyers can leave reviews regarding the accuracy of product descriptions which are then turned into coarse aggregate measures of past behavior. Our results imply that, if done correctly, this kind of aggregation can significantly increase trade. Furthermore, whether or not these platforms manage their ratings in the manner above, they do have the power to implement such a rating system.³

As a further example, consider the credibility of financial advice. A common problem is that investment brokerages (sender) have incentives to ‘oversell’ products to their clients (receivers) and are far too optimistic in their recommendations [Dugar and Nathan, 1995, Lin and McNichols, 1998, Michaely and Womack, 1999, Krigman et al., 2001, Hong and Kubik, 2003]. In 2002 the National Association of Security Dealers (NASD), a financial industry self-regulating body, imposed a new rule forcing brokerages to disclose the aggregate distribution of their recommendations. Barber et al. [2006] and Kadan et al. [2009] show that on implementation of the rule ‘Buy’ recommendations fell from 60% of total recommendations to 42%. Furthermore, prices became more responsive to ‘Buy’ calls, suggesting clients found the new reports more persuasive. We show that this is broadly in line with the predictions of our model. Finally, we use our analysis to suggest some potential improvements to this legislation.

Our third set of results concern the analysis of optimal information structures in the absence of either the coin and cup or a review aggregator. We characterize how marginal incentives to persuade differ with and without commitment. We then provide both necessary conditions and sufficient conditions for optimal information structures to be identical in the long-run and commitment cases. These conditions are closely related to the concept of a partitional information structure. Our analysis can be applied to understanding how long-run persuasion affects optimal information provision in many persuasion, cheap talk and testing problems: such as, KG, Gentzkow and Kamenica [2014, 2016a] Alonso and Câmara [2016a,b] Brocas and Carrillo [2007], Perez-Richet [2014], Che and Hörner [2015], Kolotilin [2015], Kolotilin et al., Tamura [2016], Ely [2017] Gill and SgROI [2008, 2012], Rhodes and Wilson [2017]; Chakraborty and Harbaugh [2010], Che

²Long-run players are capable of taking actions which are not best-responses at the stage game. This provides them with an ability to make transfers between periods so that they can commit to a review strategy.

³Of course, such platforms have other functions and face other constraints absent in the model examined in this paper that might complicate matters.

et al. [2013], Lipnowski and Ravid [2017].

Underlying all of our results are two technical contributions: first, we needed to show that even in the absence of commitment it is without loss to use the belief based approach of KG (Lemma 1). Second, we needed to show that the optimal information structure under commitment generically involves no two messages with the same payoff (Lemma 2). Lemma 1 allows us to reduce the sender's strategies down to a choice over information structures that map the sender's payoff relevant information into Bayes plausible distributions over the receiver's posterior beliefs at each history. We can then use the tools developed by Fudenberg, Kreps, and Maskin [1990] to characterize the optimal information structure in long-run persuasion as the solution to a static optimization problem. We use this to directly compare long-run persuasion to the cases of commitment and static cheap talk. Lemma 2 simplifies the analysis by allowing us to write off the non-generic set of cases in which static cheap talk on subsets of the state space can do as well as persuasion with commitment.

The most relevant literature to our paper is that on persuasion. Our main contribution in this area is to examine how and when the desire for credibility in a dynamic setting can micro-found the commitment assumption introduced in Kamenica and Gentzkow [2011]. We also characterize how optimal information structures differ when commitment is not exogenously imposed. These findings apply to the extensive persuasion literature discussed above, but also to the problems considered in Aumann and Maschler [1995], Rayo and Segal, 2010, Taneva, 2016 and Bergemann and Morris [2016]. While Kremer et al. [2014], Ely [2017], Bizzotto et al. [2017] consider dynamic persuasion problems, they do so with commitment at the stage game. Our paper is the first to provide such foundations for commitment.

Two papers, Perez-Richet [2014] and Piermont [2016], partially relax the degree of commitment in persuasion settings. They consider a sender who first privately learns the true distribution of states, and only then commits to an information structure for the rest of the game. However, we ground persuasion in pure cheap talk communication at every period of the game - never granting exogenous commitment at any point of the game.

Our approach of garbling histories to improve communication relates to the applied literature on online reputation and feedback systems.⁴ The closest papers to ours are Dellarocas [2005] and Doraszelski and Escobar [2012]. Both papers study forms of feedback in settings with both long and short lived players. They find that the payoff sets identified in Fudenberg and Levine [1994] are strict and achievable bounds on what can be implemented by their feedback systems. In stark contrast, our review aggregator can implement payoff vectors that strictly exceed these bounds - attaining payoffs arbitrarily close to those available under commitment. A second important contribution of this paper then, is to develop mechanisms that allow us to implement Pareto improvements

⁴See Tadelis [2016] for an excellent review.

relative to those identified in the literature thus far.⁵

Our paper has some relation to recent work on repeated Bayesian games with communication. For example, Margaria and Smolin [2015] consider long-run receivers in repeated cheap talk games. Hörner, Takahashi, and Vieille [2015] show that, when all players are long-run, it is without loss to consider only truthful equilibria for characterizing attainable payoff vectors. However, our sender cannot threaten the *short-lived* receivers, which creates a role for persuasion in our setting not present in these prior papers. Finally, Jullien and Park [2014] examine a game where communication improves the market’s ability to learn about sellers’ underlying persistent type. By contrast, the setting in our paper allows us to speak to different questions: we can characterise the limits of persuasion in a very general fashion as well as proffer mechanisms for increasing its efficiency.⁶ Also, unlike much of the repeated games literature, we focus on how to implement particular payoffs with specific strategies and information structures.

The rest of this paper is structured as follows: Section 1 provides an extended example that illustrates our two main results; in Section 2 we describe the model; in Section 3 we analyze the equilibria with perfect monitoring as well as the coin and cup; in Section 4 we analyze the review aggregator mechanism and two applications; in Section 5 we study how long-run incentives affect the sender’s optimal choice of information structure in the absence of mechanisms such as the coin and cup or review aggregator. Section 7 concludes. All formal proofs are left to the Appendix.

1 An Example

Each period t a long-run seller (with discount rate δ) tries to convince a new short-run customer (discount rate 0) to buy a product of quality $\theta_t \in \{1, 3\}$ at a fixed price, $q = 2$.⁷ The seller and customer have a common prior belief $\mu_0 = 1/3$ that the quality of each good is high ($\theta_t = 3$), where quality is i.i.d. across periods. The seller privately learns whether the quality of the good is high or low ($\theta_t = 1$) and sends a cheap talk message to the customer, who then decides whether to buy the good or not. Both the seller and the customer get a payoff of 0 if the customer does not buy. If the customer buys then the seller gets a payoff of 1; the customer gets $\theta_t - q$, and leaves a public review $\omega_t \in \{1, 3\}$. This review is informative but potentially imperfect: $Pr(\omega_t = \theta_t | \theta_t) = p \in (0.5, 1]$. When $p = 1$ the review always matches quality, $\omega_t = \theta_t$, and we say monitoring is perfect;

⁵In order to say more about persuasion, we have focused on the consequences of review aggregation for optimal communication. We expect that our results can be extended to more general environments.

⁶A talk by Laurent Mathevet on January 31st 2017 examined reputational types in our framework, based upon on going work with David Pearce and Ennio Stacchetti. They find persuasion can sometimes do well in the long-persuasion game with perfect monitoring when people believe Kamenica and Gentzkow [2011] types may exist.

⁷We ask the reader to delay worries about fixed prices as we justify this in Subsection 4.1.

otherwise, we say monitoring is imperfect. Finally, the public history is the sequence of (i) seller's past messages; and (ii) (potentially noisy) signals about the quality of goods.

We will first consider the stage game. The customer will only buy the good if she believes $\Pr(\theta_t = 3) \geq 0.5$. The solid line in Figure 1 then shows the seller's payoffs, $v(\cdot)$, as a function of the customer's posterior probability, μ_t , that the good is of high quality:

$$v(\mu_t) = \begin{cases} 1 & \text{if } \mu \geq 0.5 \\ 0 & \text{if } \mu < 0.5 \end{cases}.$$

The customer's expected payoff as a function of her posterior belief is given by $u(\mu) = \max\{2(\mu - 0.5), 0\}$. First, suppose that the seller can commit to recommending 'Buy' or 'Don't Buy' as a (potentially stochastic) function of quality. As KG show, the seller's problem can be thought of as a choice of any pair of message conditional posteriors satisfying the law of total probabilities

$$\Pr(\theta_t = 3) = \Pr(\theta_t = 3|Buy)\Pr(Buy) + \Pr(\theta_t = 3|Don't Buy)(1 - \Pr(Buy)), \quad (1)$$

Hence, we can think of a feasible information structure as any choice of posteriors $\mu_D := \Pr(G|Don't Buy)$ and $\mu_B := \Pr(G|Buy)$ around μ_0 . The customer's preferred information structure is truth-telling, $(\mu_D, \mu_B) = (0, 1)$. For the seller, this information structure pays out for high quality goods only. The payoff is represented in Figure 1 by the weighted average of $v(0)$ and $v(1)$, which lies above μ_0 .




Figure 1: Seller's payoffs given Customer's posteriors

However, the seller can get a better stage payoff with the policy $(\mu_D, \mu_B) = (0, 0.5)$ which concavifies the value function as in KG. This information structure maximizes the unconditional probability of sending an incentive compatible *Buy* message to the seller; it is achieved by sending a *Buy* message whenever the quality is high, and half the time when the quality is low; the seller then sells the good with an ex-ante probability of $\frac{2}{3}$ (Figure 1). While this information structure gives the seller his highest expected stage

payoff it gives the customer zero expected surplus, irrespective of the message.

Of course, without commitment, there is no equilibrium information structure of the stage game in which the customer follows the seller's advice. All equilibria are payoff equivalent to babbling – where the seller always reports *Buy* and the customers ignore him. However, we can do better in the long-run setting where the seller maximizes his expected lifetime discounted utility,

$$V_0 = \sum_{t=0}^{\infty} \delta^t v(\mu^t),$$

and customers form their posteriors based on both the message they receive *and* the public history. This is because a period of babbling can now be used to punish the seller whenever there is a bad review, $\omega_t = 1$, and hence can help generate credibility at the stage game.

To be concrete, consider a putative truth-telling equilibrium where a period of babbling reduces the seller's discounted equilibrium continuation payoffs by $1/p$ when he reports *Buy* and subsequently gets a bad review. The probability of being punished for sending *Buy* when $\theta = 1$ is p . Consequently, the expected punishment for this is just equal to $1 -$ the stage payoff from selling the good. The seller then is indifferent between messages when quality is low, strictly prefers to send *Buy* when the quality is high, and so truth-telling is incentive compatible.

Notice two things here: first, when the seller sends *Buy* and $\theta_t = 3$ there is an expected on path punishment of $(1-p)/p$ from a bad review being left incorrectly; and second, $1/p$ is the most efficient punishment that satisfies incentive compatibility. The value of this equilibrium to the seller (and receiver) at any stage where the seller is not being punished with babbling is

$$V^{T'} = \frac{1}{3} \left(1 - \frac{1-p}{p} \right) = \frac{2p-1}{3p} > 0. \quad (2)$$

As this payoff is positive, babbling for a finite period of time can induce a punishment of $1/p$ for sufficiently patient sellers and hence enforce the equilibrium. Further, as a punishment of $1/p$ makes the seller indifferent between messages when $\theta_t = 1$, it is easy to see that we can use this punishment to support any equilibrium information structure such that $\mu_D = 0$ and $\mu_B \in [0.5, 1]$. However, all of these information structures only provide the seller with the *same* discounted average payoffs. Indeed, the seller only picks up positive payoffs (net of punishments) when the quality of the good is high. The average payoffs of customers however is increasing in μ_B . The available equilibrium payoffs are described in figure 2a below:

The full-triangle *BTP* describes the stage payoff pairs available at any period. The

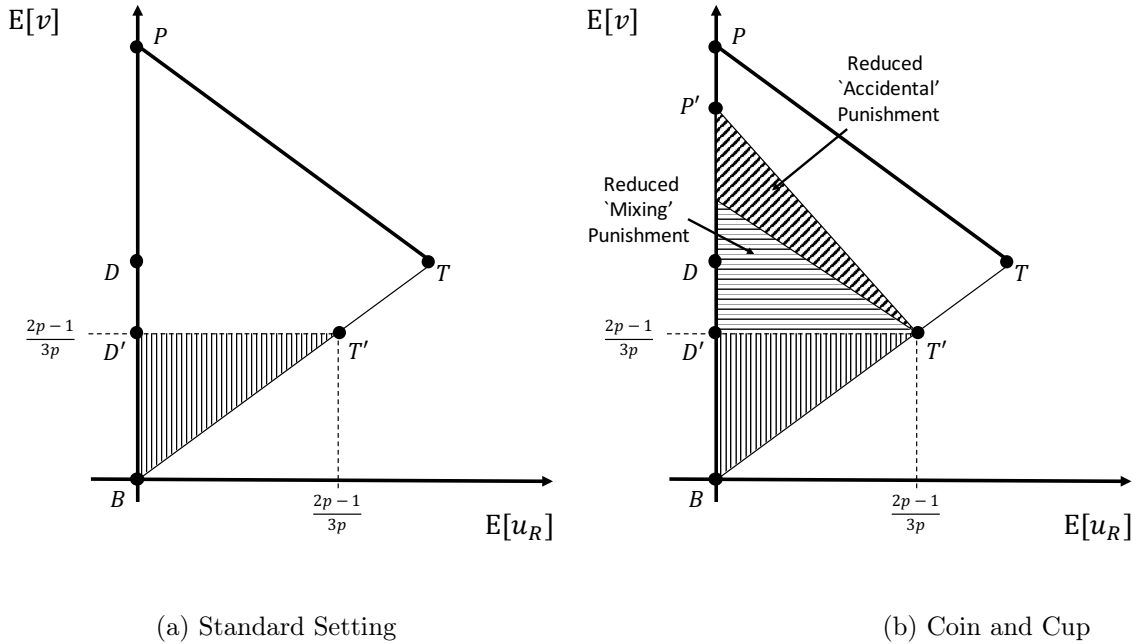


Figure 2: Average Payoff Sets

lower triangle $BT'D'$ describes the average payoff pairs available in equilibrium. The sole efficient information structure here is truth-telling – all other policies can only (weakly) reduce the seller’s discounted average payoffs, and strictly harm receivers. Notice that, as monitoring becomes perfect, the payoff sets tend to $BT D$ - a strict subset of what can be achieved under commitment (the triangle BTP).

A Coin and Cup

Now consider a slight change to the stage game where we augment the history with temporary asymmetric information. At the beginning of each period the seller shakes a coin in a cup, places it on the table and peeks under the cup to see whether the coin came up heads or tails. The customer observes the seller do all this, but does not see the coin. The cup, with the coin still under it, is left on the table. Then the seller learns the quality of the good, sends a report, and the customer makes her decision, as before. After the decision the customer lifts the cup, observes the coin, and now records whether the coin was heads or tails with her review. All customers now observe the history of messages, reviews and coin flips.

We can now look for equilibrium information structures that map both the quality and the value of the coin into Bayes plausible posteriors. In particular, the seller can replicate $(\mu_D, \mu_B) = (0, 0.5)$ by reporting *Don't Buy* if and only if $\theta = 1$ and the coin comes up tails. This is enforced with a punishment of $1/p$ whenever the seller leaves a bad review and the coin is recorded as tails.⁸ Relative to the same information structure without the coin and cup, this has two benefits. First, it allows for a pure strategy that is *ex-ante*

⁸Alternatively one might think of this as a garbling of reviews such that when the coin comes up heads the seller knows the review will be good whatever message he sends.

stochastic from the perspective of the customer yet is verifiable *ex-post* as deterministic (imperfect monitoring issues aside). This allows the seller to get an extra payoff of 1 when the good is low quality and the coin is heads - an expected gain of half. Second, it reduces “accidental” punishment where $\theta_t = 3$ but the customer leaves a bad review; this has an expected value to the seller of $(1-p)/2p$. These two benefits are described in figure 2b above.⁹

If we use a continuous random variable for “the coin” then we can attain the payoff set of $BT'P'$ in equilibrium. As ex-post monitoring of the seller’s payoff relevant information set becomes perfect we move towards BTP , achieving it at $p = 1$. That is, under perfect monitoring, a coin and cup allows the seller to costlessly commit to the optimal commitment information structure without the need for third party verification, contracts, or transfers. While the coin and cup costlessly substitutes for commitment when $p = 1$ it does not when $p < 1$ and we can still be far from the frontier PT . This raises the question of whether a larger set of equilibrium payoffs can be achieved by augmenting, or garbling, the public history further – a question to which we now turn.

A Review Aggregator

In general, if monitoring is imperfect but *all* players are sufficiently patient we know that almost any average payoffs are sustainable, subject to some identifiability conditions. Hence, by conditioning punishments on a large number of observations, might short-lived customers strip out the noise in the individual reviews? The answer is no - this follows fairly straightforwardly from Fudenberg and Levine [1994]. However, it turns out that if we introduce a “Review Aggregator” who garbles the history we can run a statistical test and get arbitrarily close to commitment payoffs and information structures.

The review aggregator provides the customers with a binary aggregate rating, good or bad, of the seller’s past behavior. A good rating lasts for T individual reviews. If the number of individual bad reviews are below some critical threshold, $X < T$, then the seller “passes the test” and keeps his good rating for a further T individual reviews. If the seller gets X or more bad individual reviews then he “fails the test” and his aggregate rating becomes bad for a finite number of periods Z . The customers only observe the current aggregate rating – in particular they do not see the full history of play and they are kept uncertain about how many reviews there are before the next test. The seller does not see the individual reviews either, but he knows how many reviews there will be before he is next tested. As before, the seller sends messages directly to the customers

⁹The first benefit is comparable to the gains that come from being able to observe mixed strategies and identified in Fudenberg, Kreps, and Maskin [1990]. However, the second benefit, as far as we have been able to ascertain has not been identified. The coin and cup can be seen as a device for replacing the random shocks considered in Aumann [1961] and Harsanyi [1973] with something that is ex-post observable. By doing this we are able to still think about observable mixed strategies in an environment where information sets (or actions) are imperfectly observed.

who then make a decision based on the aggregate rating and the message.

Suppose the seller is arbitrarily patient and the customers always obey advice when there is a good aggregate rating. A straightforward extension of Radner [1985] shows that with the correct choice of T , X and Z , we can achieve almost any *average* information structure such that $(\mu_D = 0, \mu_B \in (0.5, 1))$ with negligible on path punishments. This implies further, that we get *average* payoffs arbitrarily close to the Pareto frontier PT . While the average information structure may be $(\mu_D = 0, \mu_B \in (0.5, 1))$ the particular information structure that the seller chooses will depend on the number and quality of goods sold since the last test. In particular, there are histories at which the seller will choose to always send *Buy*. If at that history the customer could see the number of goods sold since the last test, or even all the reviews, she could infer that the seller will send *Buy* regardless of the state. However, as the customer knows neither the number nor quality of goods sold since the last test she just faces the average information structure: $(\mu_D = 0, \mu_B \in (0.5, 1))$. Consequently, it is incentive compatible for each customer to follow the seller's advice.

The job of the review aggregator here, is providing some record of the seller's behavior while keeping each customer uncertain about the precise information structure they are facing at t . Without the aggregator, following advice is only ever incentive compatible for the *short-run* customer at t if she believes $\mu_B \in [0.5, 1]$ at that precise history. This is only possible if the sender is punished for a bad review at t by at least $1/p$. Hence, any positive stage payoff for the seller must be associated with on-path punishments holding us down to the payoffs in (2).

In order to effectively use a statistical test to reduce on path punishments the seller has to be allowed to deterministically report *Buy* at some histories while still having his advice obeyed by the customer. This is not incentive compatible though when the customers see the complete history, and so we need the review aggregator. What is striking here, is how little we need to relax these constraints to achieve payoffs arbitrarily close to those available under commitment. For instance, if one were to consider enforcing truth-telling, the seller would only lie deterministically on an arbitrarily small proportion of occasions. Yet, if we don't allow this small number of deterministic lies the tool of the statistical test comes tumbling down.

2 The Model

A sender S ('he') plays the following stage game against an infinite sequence of receivers R (each 'she').

Stage Game

Each period, a one-period receiver R_t must take an action a_t from a compact set A . Her payoffs from action a_t depend on an unknown state of the world, $(\theta_t, \omega_t) \in \Theta \times \Omega = \{\theta^1, \theta^2, \dots, \theta^N\} \times \{\omega^1, \omega^2, \dots, \omega^{N_\omega}\}$. Her payoffs are given by the utility function $u_R(a_t, \theta_t, \omega_t)$. Each θ_t is drawn independently across time, from a prior distribution represented by the vector $\mu_0 \in \Delta\Theta$.¹⁰ Conditional on θ_t , ω_t follows the distribution $g(\omega_t | \theta_t) > 0$, where ω_t is also independent of $\theta_\tau, \omega_\tau, \tau \neq t$.¹¹ However, note that we allow for θ_t and ω_t to be contemporaneously correlated.

At each t , a receiver R_t arrives ex-ante uninformed about (θ_t, ω_t) and leaves the game at the end of the period. At the beginning of each period, an infinitely-lived sender S privately observes the realization, θ_t . Before R_t takes an action, S sends a message m_t from some set, M . With some notational abuse, we sometimes refer to the cardinality of the message space by M . Within a period, the sender only cares about the action taken by R_t and has stage utility $u_S(a_t)$.

Within period, the timing of this *static cheap talk game* is as follows:

1. θ_t, ω_t are drawn respectively from distributions $\mu_0, g(\omega_t | \theta_t)$. S privately observes θ_t .
2. S sends a message $m_t \in M$ (possibly random) to R_t .
3. After observation of m_t , R_t chooses an action $a_t \in A$.
4. After taking action a_t , ω_t is observed by all players.

We interpret θ_t as the sender's private information relevant for receiver's decision problem, and ω_t as ex post feedback that the receiver learns only after taking action. Written this way, the model is flexible enough to describe interactions where (i) the sender is better informed about receiver's preferences, $u_R(a, \theta)$; (ii) receiver's final information of her own preferences after taking action will be better than the sender's, $u_R(a, \omega)$; and (iii) neither the sender nor receiver information can be ranked as superior for decision-making, $u_R(a, \theta, \omega)$.¹²

After receiving message m_t , the receiver R_t forms her posterior belief μ_t and chooses her action $a_t(\mu_t)$ to maximize $\mathbb{E}[u_R(a_t, \theta_t, \omega_t) | m_t]$. We often refer directly to the sender's equilibrium period- t stage payoff, as a function of the receiver's posterior:

$$v(\mu_t) := u_S(a(\mu_t)).$$

¹⁰Following standard notation, we use ΔX to denote the simplex over set X .

¹¹Note in particular that we assume the conditional distribution of ω_t has full support, for all θ_t . While this assumption is not strictly necessary for the main results, it significantly aids exposition of our results related to the review aggregator in Section 4.

¹²These functions can be used to think respectively about (i) uncertainty over a good's quality; (ii) experience goods with heterogeneous preferences; or (iii) asset and/or matching markets.

As in KG, we assume that whenever R_t 's posterior belief leaves her indifferent between two or more actions, she chooses the one S prefers. This ensures that $v(\mu_t)$ is an upper semi-continuous function. We refer to this stage game by Γ_t . As the stage game is a standard cheap talk game, there always exists a babbling equilibrium in which S 's messages are completely uninformative.

We will contrast this game with a *static information design* problem, in which S can *commit* in advance to a (mixed) reporting strategy before learning θ_t . In the persuasion game, the timing and available actions are as follows:

1. S chooses an *experiment*: a message space M , and a random mapping $\hat{s} : \Theta \rightarrow \Delta M$.
2. θ_t is privately drawn from distribution μ_0 . Conditional on θ_t , $m_t \in M$ is drawn from s_0 .
3. R_t observes m_t and chooses an action $a_t \in A$.

In the static information design problem, S commits (before observing θ_t) to an *experiment* (a message space M , and a garbling \hat{s} of θ_t). The key distinguishing feature of an experiment is that S can commit to any (potentially stochastic) policy.

Define $\hat{v}(\mu)$ as the smallest concave function that is everywhere weakly greater than $v(\mu)$. That is,

$$\hat{v}(\mu) := \sup \{ \nu : \nu \in co(v) \}$$

where $co(v)$ denotes the convex hull of the graph of v . KG show that S 's optimal payoff via information design is exactly $\hat{v}(\mu_0)$, the *optimal commitment payoffs*.

By definition, $\hat{v}(\mu_0) \geq v(\mu_0)$. If $\hat{v}(\mu_0) = v(\mu_0)$, then S 's optimal payoff can be achieved by sending no information to R_t , or by a babbling equilibrium of the cheap talk game. To ensure that persuasion is a useful tool for the sender, we assume in the rest of the paper that v, μ_0 are such that

$$\hat{v}(\mu_0) > v(\mu_0).$$

The Repeated Game

The stage game Γ_t is played again against a new receiver each period $t = 0, 1, 2, \dots$. We refer to this long-run persuasion game by Γ^∞ . At each period t and public history $\phi_t = (m_\tau, a_\tau, \omega_\tau)_{\tau=0}^{t-1}$, the sender and receiver R_τ observe ϕ_t (the sender also observes the private history $\theta^t = (\theta_\tau)_{\tau=0}^t$) and play game Γ_t . The sender's discounted payoff from a sequence of receiver actions $a = (a_1, a_2, \dots)$ is

$$\sum_{t=0}^{\infty} \delta^t u_S(a_t).$$

Let the set of all period- t histories be Φ_t . At period t , let the map $s_t : \Phi_t \times \Theta^t \rightarrow \Delta M$ express a history and state dependent probability distribution over the sender's messages. A strategy for the sender is a collection $s = (s_t)_{t=0}^\infty$. Similarly, let a mixed strategy for receiver R_t be a map $\rho_t : \Phi_t \times M \rightarrow \Delta A$.

We use the term equilibrium to refer to Perfect Bayesian Equilibria of the above game. An equilibrium specifies: a strategy s for the sender; strategies $\rho = (\rho_t)_{t=0}^\infty$ for each R_t ; and history-dependent posterior belief functions $\{\mu_t\}_{t=0}^\infty$, where $\mu_t \in \Delta\Theta$ is an N -dimensional vector, such that:

1. Given the receivers' strategies and history (ϕ_t, θ_t) , s maximizes the sender's expected discounted payoff $\mathbb{E}[\sum_{\tau=t}^\infty \delta^\tau u_S(a_t) \mid \phi_t, \theta_t; \rho]$.
2. Given the sender's strategy, ρ_t maximizes R_t 's expected payoff, $\sum_{i=1}^N \mu_t^i \cdot u_R(a, \theta_t^i)$.
3. Where possible, the receiver's posterior beliefs $\mu_t = (\mu_t^1, \dots, \mu_t^N)$ satisfy

$$\mu_t^i = \Pr(\theta_t = \theta_t^i \mid \phi_t, m_t; s).$$

Regardless of play in rounds $\tau < t$, condition 3 above ensures messages in the support of the sender's strategy at any history ϕ^t must be consistent with the prior belief over θ_t . However, for off-path messages chosen at time t , equilibrium places no restrictions on receiver R_t 's beliefs.

A Direct Equilibrium

Building on an insight from KG, the notion of equilibrium in the long-run persuasion game can be cast entirely in terms of history-dependent lotteries over beliefs, μ_t . Define the stage game $\hat{\Gamma}_t$ as follows: the sender sends message $\tilde{\mu}_t$ from the set of possible posterior beliefs that R_t may hold, $\Delta\Theta$, and is elsewhere the same as Γ_t . The long-run persuasion game, $\hat{\Gamma}^\infty$, is analogously defined. In such an environment, histories are now vectors of the form $h_t = (\tilde{\mu}_\tau, a_\tau, \omega_\tau)_{\tau=0}^{t-1}$, the set of all period- t histories H_t , and (behavioural) strategies $\sigma = (\sigma_\tau(h_\tau, \theta^\tau))_{\tau=0}^\infty$, where each $\sigma_t : H_t \times \Theta^t \rightarrow \Delta M$, and $\rho_t : \Phi_t \times M \rightarrow \Delta A$ for S, R_t respectively. We denote the set of all strategies for S by Σ .

Let the function $\mu_t(h_t, \tilde{\mu}_t)$ specify R_t 's beliefs given history h_t and message $\tilde{\mu}_t$. We define the concepts of information structure as follows:

Definition 1. An *information structure* is a lottery $\lambda = (\lambda_j, \mu_j)_{j=1}^M \subset \Delta(\Delta\Theta)$ over a set of $M \leq N$ posteriors $\{\mu_j\}_{j=1}^M$ where $\lambda_j = \Pr(\mu = \mu_j)$.

An information structure is *Bayes plausible* if $\mu_0 = \sum_{j=1}^M \lambda_j \mu_{t,j}$. We refer to the set of Bayes plausible information structures as $\Lambda(\mu_0) \subset \Delta(\Delta\Theta)$. We can now define a *direct equilibrium* of the long-run persuasion game as follows:

1. Given the receivers' belief functions $\mu_t(h_t, \tilde{\mu}_t)$, $\tilde{\mu}_t \in \text{supp}(\sigma_t(h_t, \theta_t))$ maximizes the sender's expected discounted payoff

$$V_t(h_t, \theta_t) = v(\mu_t(h_t, \tilde{\mu}_t)) + \delta \mathbb{E}[V_{t+1}((h_t, \tilde{\mu}_t, \theta_t), \theta_{t+1})] \quad (3)$$

where V_t is the sender's continuation payoff at history (h_t, θ_t) .

2. The receiver has obedient beliefs: $\mu_t(h_t, \tilde{\mu}_t) = \tilde{\mu}_t$ for all $\tilde{\mu}_t \in \cup_{\theta_t \in \Theta} \text{supp}(\sigma_t(h_t, \theta_t))$.
3. The information structure, λ_t , is Bayes plausible and $M_t \leq N$ for all h_t .

V_t is simply the sum of S 's discounted payoffs from equilibrium play at history (h_t, θ_t) onwards. In any equilibrium, S must maximize (3) at all histories of the game tree, given $\mu_\tau(h_\tau, \tilde{\mu}_\tau)$, $\tau \geq t$. Moreover, a *direct equilibrium* requires that (i) R_t 's beliefs conform to the recommendation made by S , for any $\tilde{\mu}_t$ on the equilibrium path, (ii) at any history, S 's mixed strategy over messages can be 'averaged back' to the prior μ_0 . Finally, given a belief μ_t , recall that the optimal equilibrium behavior of the receiver is implicitly included in the sender's value function, $v(\mu_t(h_t, \tilde{\mu}_t))$.

The following Lemma establishes is without loss to restrict attention to *direct equilibria* of game Γ^∞ :

Lemma 1. *For any equilibrium of game Γ^∞ , there is a direct equilibrium of game $\hat{\Gamma}^\infty$ that induces the same distribution over receivers' actions, for each state θ_t and history h_t on the equilibrium path.*

In order to prove Lemma 1 we cannot directly employ the approach of KG as we do not have full commitment. Instead, we show that we can map any non-direct strategy into a direct strategy such that the equilibrium conditions 1 to 3 are satisfied.

3 The Limits of Long-Run Persuasion

In this section, we establish necessary and sufficient conditions under which long-run persuasion and persuasion under commitment have the same optimal payoffs. As a preliminary we show that for any typical persuasion problem under commitment, it is without loss to restrict attention to cases where the optimum is achieved only by information structures involving posteriors that have different values:

Lemma 2. *Generically, the concave envelope at μ_0 , $\hat{v}(\mu_0)$, is supported by distinct posteriors $\{\mu_j^*\}_{j=1}^M$ such that $v(\mu_j^*) \neq v(\mu_l^*)$ for any $j \neq l$.*

Lemma 2 establishes that the optimal commitment information structure, λ^* , typically involves a lottery over posteriors over which the sender has strict preferences. Formally, we use a convex polyhedral approximation to the subgraph of \hat{v} to establish that the set

of value functions for which the above condition holds is open and dense. One important implication of this is:

Corollary 1. *Generically, optimal commitment payoffs cannot be achieved in a static cheap talk game.*

Since cheap talk requires the sender to be indifferent across all messages, Lemma 2 tells us that such equilibria are generically not optimal for the sender. Hence, generically there is a strict wedge between the payoffs achievable with and without commitment, respectively. This gap provides the upper and lower bounds on what can be achieved by long-run persuasion.

Given Lemma 2, we restrict attention to the following generic class of problems:

Assumption 1. *The concave envelope at μ_0 , $\hat{v}(\mu_0)$, is not supported by distinct posteriors $\{\mu_j^*\}_{j=1}^M$ such that $v(\mu_j^*) = v(\mu_k^*)$ for any $j \neq k$.*

3.1 Perfect Monitoring

Here we analyze the case where S's signal always matches the feedback i.e. $\theta_t = \omega_t$ for all t . We do this because the perfect monitoring case represents an upper bound on the value of repetition to persuasion. We will show that even in this setting, long-run persuasion cannot typically achieve optimal commitment payoffs.

To analyze the problem, we apply Lemma 1 and the tools of Fudenberg, Kreps, and Maskin [1990] to reduce the search for an optimal information structure in long-run persuasion to a static optimization problem that is directly comparable to both persuasion with commitment and to static cheap talk. As will be shown, the optimal Bayes plausible information structure for each type of problem is given by:

$$\lambda^{KG} \in \arg \max_{\lambda \in \Lambda(\mu_0)} \sum_{\Theta} \lambda_j v(\mu_j); \quad (\text{Commitment})$$

$$\lambda^{LP} \in \arg \max_{\lambda \in \Lambda(\mu_0)} \sum_{\Theta} \mu_0^i \min \{v(\mu) : \mu \in M^i\}; \quad (\text{Long-Run Persuasion})$$

$$\lambda^{CT} \in \arg \max_{\lambda \in \Lambda(\mu_0)} \sum_{\Theta} \mu_0^i \min \{v(\mu) : \mu \in M\}; \quad (\text{Static Cheap Talk})$$

where $M^i := \{\mu_j \in \text{supp}(\lambda) : \mu_j^i > 0\}$ is the subset of messages sent with positive probability in state θ_i . To see where these equations come from, consider first the long-run persuasion problem of maximizing S's period-0 discounted utility across all possible equi-

libria of the long-run persuasion game:

$$\begin{aligned}
& \max_{\sigma \in \Sigma} \mathbb{E}_\theta [V_0(\theta_0)], \text{ such that for all } h_t \in H_t : \\
& v(\mu_t) + \delta \mathbb{E}[V_{t+1}((h_t, \mu_t, \theta_t), \theta_{t+1})] \geq v(\mu'_t) + \delta \mathbb{E}[V_{t+1}((h_t, \mu'_t, \theta_t), \theta_{t+1})]; \\
& \text{for all } \mu_t \in \text{supp}(\sigma_t(h_t, \theta_t)) \text{ and } \mu'_t \notin \text{supp}(\sigma_t(h_t, \theta_t)); \\
& \mu_0 = \sum_{j=1}^M \lambda_j \mu_{t,j}.
\end{aligned} \tag{4}$$

Problem (4) involves choosing a strategy profile $\sigma = (\sigma_1, \sigma_2(h_2), \dots)$ for S that maximizes his present discounted utility, such that at each history h_t : (i) the choice of message is optimal for S given R_t 's beliefs; (ii) beliefs satisfy Bayes plausibility. There is a subtle difference between problem (4) and the description of equilibrium. In equilibrium, S need only maximize his choice of $\tilde{\mu}_t$ at each history h_t , subject to R_t 's beliefs. In problem (4), when we choose a strategy σ' , we are also able to vary R_t 's beliefs following any message sent, so long as they conform to equilibrium restrictions. In addition, the choice of strategy must be optimal for S , given the receivers' beliefs.

To characterize the solution to problem (4), we introduce some notation. Let $v_i(\lambda) := \min \{v(\mu) : \mu \in \text{supp}(\lambda), \mu^i > 0\}$ be the minimum payoff to S among all posteriors μ that (i) are in the support of distribution $\lambda \in \Delta(\Delta\Theta)$ and (ii) occur with strictly positive probability conditional on state θ_t^i (under λ). Then we have:

Proposition 1. *The sender's discounted average continuation value from any long-run persuasion game is bounded above by*

$$(1 - \delta) \mathbb{E}_\theta [V_0(\theta_0)] \leq \max_{\lambda \in \Lambda(\mu_0)} \sum \mu_0^i v_i(\lambda) \tag{5}$$

There exists $\underline{\delta}$ such that $\forall 1 > \delta \geq \underline{\delta}$, this upper bound can be attained at some equilibrium.

Equation (5) bounds S 's payoffs above by the best expected statewise-minimal payoff among all lotteries of posteriors $\lambda \in \Lambda(\mu_0)$. Moreover, Proposition 1 says this is attainable for sufficiently patient S . The intuition is as follows: if in some state, θ_t , S wants to mix between messages with different stage payoffs then he must be indifferent. This indifference requires that all messages but the worst induce an on-path punishment that wipe out the stage gains of that message - pinning down the upper-bound. We can sustain this equilibrium for sufficiently patient S by threat of the worst cheap talk equilibrium, as in the earlier example. The logic here is very similar to that in Fudenberg, Kreps, and Maskin [1990]. However, by restricting analysis to direct equilibria of a communication game we are able to say more about the form of a sender's objective function and, ultimately, his behavior in terms of optimal information structures.

One might worry that the above discussion does not mention enforceability constraints - what if the worst cheap talk equilibrium is too attractive? However, at the optimum it necessarily follows that this enforceability constraint is slack. In the worst case $\lambda^{LP} = \lambda^{CT}$

which is itself enforceable.

The nature of the sender’s problem in long-run persuasion now demonstrated, we can turn to how this relates to the problem under commitment and in static cheap talk. Equation (Commitment) trivially represents the optimal commitment problem in KG. However, formally, equation (Static Cheap Talk) is missing a constraint - cheap talk requires indifference across all messages. Lipnowski and Ravid [2017] show that this constraint is slack.¹³

These three equations will help us compare both the average payoffs and optimal information structures under different forms of persuasion. In this section we will be using them to study only the payoffs and will apply them to the analysis of optimal information structures in Section 5. That said, we need to introduce a particular class of information structure before using these results to establish necessary and sufficient conditions for long-run persuasion to achieve the same payoffs as the commitment benchmark:

Definition 2. An *information structure* is *partitional* if and only if it satisfies:

$$\begin{aligned} \text{supp}(\mu_j) \cap \text{supp}(\mu_k) &= \emptyset \quad \forall j \neq k \\ \lambda_j &= \sum_{i \in \text{supp}(\mu_j)} \mu_0^i \quad \forall j \\ \mu_j^i &= \frac{\mu_0^i}{\lambda_j} \quad \forall i \in \text{supp}(\mu_j) \end{aligned}$$

Such an information structure is implemented by deterministically mapping states into posteriors – this can be many to one, but not one to many. With this definition in hand we can establish the following:

Theorem 1. *Under perfect monitoring there exists a $\underline{\delta}$ such that optimal commitment payoffs are attainable in the long-run persuasion game if and only if the optimal information structure under commitment is partitional, for all $\delta \geq \underline{\delta}$.*

First, note that an immediate implication is that persuasion with commitment weakly dominates long-run persuasion. For any state-contingent lottery the sender gets the state-wise average value of the messages under commitment but the state-wise minimum value with long-run persuasion. Hence, any two lotteries can only induce the same average payoffs across the two problems when the state-wise average is equal to the state-wise

¹³Equation (Static Cheap Talk) tells us that any information structure with different value messages is still pinned down to the lowest message. Hence, the best the sender can do is the quasi-concave envelope which is the optimal cheap talk payoff.

It’s worth noting further, that this set can contain information structures not achievable by static cheap talk. In fact, this equation describes the set of optimal information structures for long-run persuasion with *no monitoring* i.e. θ_t independent of ω_t . In this case, the sender has to be indifferent between *all* messages - hence, on-path punishments pin him down to the best payoffs in the static cheap talk problem. This provides a stark contrast to the folk theorem in Margaria and Smolin [2015]: with short-run receivers and no monitoring we can’t do better than static cheap talk.

minimum. There are two ways this can happen: one, there is only one message sent in any state; and two, when there are multiple messages sent in some state, they have the same payoff. It then follows that long-run persuasion achieves the same payoffs as persuasion under commitment if and only if λ^{KG} satisfies one of these two conditions. The first case is simply a partitional information structure. The second case, however, is ruled out as non-generic following Lemma 2.

3.2 A Coin and Cup

We now introduce a novel ‘coin and cup’ mechanism which can always retrieve optimal commitment payoffs under perfect monitoring. A ‘coin and cup’ mechanism introduces a payoff-irrelevant variable, $c_t \sim U[0, 1]$, drawn independently of θ :

1. θ_t, c_t are drawn from their respective distributions. S privately observes **both** realizations.
2. S sends a message $m_t \in M_t$ to R_t . R_t **only** observes the message.
3. After observation of m_t , R_t chooses an action $a_t \in A_t$.
4. After taking action a_t , c_t and the state θ_t are observed by all receivers.

As the next theorem shows, the coin and cup allows the sender to achieve his full commitment payoffs:

Theorem 2. *For any payoff available to S from some experiment, there exists $\underline{\delta}$ such that the payoff is attainable with a coin and cup whenever it exceeds S ’s worst stage payoff, for $\delta \geq \underline{\delta}$.*

With a coin and cup, all information structures can be implemented as partitions of a larger ‘state space’, consisting of the unit interval and Θ .¹⁴ In this way, a deterministic strategy conditioned on both θ_t and c_t can appear mixed from the perspective of the receiver that period, but ex-post is verifiable. Hence, under perfect monitoring, the threat of punishment never need be exercised.

A useful feature of a coin and cup is that it is implementable without any specialist knowledge of the decision problem or the realization of θ . Further, we do not require S to be able to commit to a specific experimental procedure for generating a particular distribution, c_t . As the simplest possible example described in Section 1, we can do this with something as mundane as a coin and cup - or many coins and a time-lock safe. As a more high-tech example, blockchain technologies (such as that underpinning Bitcoin) support decentralized recording and updating of information among peers using

¹⁴Such a partitional mapping is identical to the definition of a signal introduced by Green and Stokey [1978] and used later in Gentzkow and Kamenica [2017].

cryptographic methods. These technologies can be used to share information in a way that cannot subsequently be tampered with, and allow for information to be withheld from some participants until pre-specified times.

In one sense, the coin and cup is a tool that substitutes for the observable mixed strategies examined in Fudenberg, Kreps, and Maskin [1990]. However, it more closely resembles the models of mixed strategies discussed in Aumann [1961] or Harsanyi [1973] where agents condition pure strategies on privately observed variables. By taking this approach we are also able to see that such a device is not only useful under perfect monitoring but also under imperfect monitoring. Indeed, as we saw earlier, it allows a larger set of feasible average payoffs than in Fudenberg and Levine [1994] who did not examine observable mixed strategies under imperfect monitoring.

4 Review Aggregation

In Section 1 we saw that under imperfect monitoring both average payoffs and information structures were strictly bounded away from the case with exogenous commitment. This result extends more generally by a straightforward application of the analysis in Fudenberg and Levine [1994]. However, we will now show that a review aggregator can overcome these constraints on payoffs. With the aggregator, we can now get arbitrarily close to the full set of payoffs available under commitment.

For the purpose of exposition, we will use the payoffs, priors, and actions of the example in Section 1. We will consider a symmetric binary review of the sender’s information set $Pr(\omega_t = \theta_t | \theta_t) = p$, where $1 > p > \frac{1}{2}$, that has no direct impact on the receivers’ stage payoffs. For a generalization of these results, we refer the reader to the Online Appendix.

The role of the review aggregator will be to garble the receivers’ observations of the sender’s past interactions in a particular way. Importantly however, it does not interfere with the direct communication between the sender and receiver. More specifically, the aggregator changes the model of Section 2 as follows.¹⁵ At time 0, the aggregator announces to the first receiver that the sender has a ‘Good’ rating, \mathcal{G} . Subsequently, the sender and current receiver play the cheap talk game as before. This is repeated for T interactions, and at each time the incoming receiver is told only that the sender’s past play means that his rating is ‘Good’. We call this block of interactions a ‘ \mathcal{G} -phase’.

Importantly, receivers observe *nothing else*. We let the aggregator randomly permute the ordering of the T receivers so that they cannot infer their position in the line from their index, t . In our context this is a simple way of capturing the uncertainty about the relationship between the calendar date and the actual number of interactions that have taken place when receivers only observe an aggregate statistic.¹⁶

¹⁵A fuller, formal description of the aggregator’s role is postponed to the Online Appendix.

¹⁶Alternatively, we could model this by having the review aggregator generate uncertainty by ran-

At the end of the first T periods, the aggregator conducts a statistical test of the sender's strategy compared to some benchmark (mixed) strategy, $\sigma^* : \Theta \rightarrow \Delta(\Delta\Theta)$, chosen in advance. This is done by comparing the empirical joint distribution of reports $(\tilde{\theta}_t)_{t=0}^{T-1}$ and subsequent feedback, $(\omega_t)_{t=0}^{T-1}$. If the sender is sufficiently likely to have played a strategy 'close to' σ^* he passes the test and another \mathcal{G} -phase begins. Otherwise, the aggregator switches to announcing to the next $\beta T \in \mathbb{N}$ receivers that the sender's rating is 'bad', \mathcal{B} .

We establish the following Theorem:

Theorem 3. *A review aggregator can achieve average payoffs arbitrarily close to those under commitment.*

Theorem 3 shows that we can now achieve the full commitment payoffs, demonstrating the value of review aggregators, such as those used by eBay, for sustaining efficient communication and persuasion. This provides a further micro-foundation for the commitment assumption in the persuasion literature. The proof of the Theorem is constructive. In any \mathcal{B} -phase, both senders and receiver play a babbling equilibrium. This provides the punishment to a sender for deviating 'too far' from σ^* on average during any \mathcal{G} -phase. The test is chosen to be strict enough that S 's optimal strategy can be bounded close enough to σ^* on average by the expected costs of falling into a \mathcal{B} -phase, but lenient enough that if S did adopt σ^* each period he would pass with a high probability. Finally, we show that the bounds on S 's optimal strategy, along with the uncertainty they face about the history of play, are enough to satisfy each receiver's incentive constraint.

Our review aggregator plays two roles in making the full set of commitment payoffs feasible. First, in its interaction with S , it uses the individual reviews ω_t to statistically monitor the sender's strategy (as in Radner [1985]). By doing this, we can reduce on path punishment, almost costlessly ensuring the sender plays close to σ^* .

The aggregator's second role, however, is novel. As we discussed above, short-run receivers cannot profitably agree to conduct these tests. By obfuscating the history of the Sender's past interactions in a \mathcal{G} -phase, the aggregator makes it much easier to satisfy the receiver's incentive constraints. Instead of requiring that the sender's strategy be credible at each and every instant in the game, a receiver now only needs to be satisfied that the sender plays close to σ^* *on average* (across T). Given this, and a lack of knowledge about where she is in the \mathcal{G} -phase, she now finds it optimal to follow sender's reports. This is crucial to achieving payoff sets that are strictly larger than those identified by Fudenberg and Levine [1994].

This slight relaxing of the Receiver's incentive constraint is powerful. By breaking the need for credibility each and every period, the sender can now obtain benefits from domising the lengths of each ' \mathcal{G} -phase' and ' \mathcal{B} -phase'. This complicates the analysis considerably and so we simplify the analysis by using this assumption. Gershkov and Szentes [2009] and Kremer et al. [2014] make similar assumptions in different contexts.

persuasion without the need for being punished. For instance, in the example of Section 1, suppose we set σ^* such that $\Pr(m = \text{Buy} \mid \theta = 3) = 1$, $\Pr(m = \text{Buy} \mid \theta = 1) = \frac{1}{2} - \varepsilon$, $\varepsilon > 0$. For ε small enough, the sender can always guarantee close to the optimal commitment payoff by playing according to σ^* each period. In doing so, he secures the optimal commitment payoff during the \mathcal{G} -phase, and he is highly likely to pass the test. Each receiver follows his recommendations because (i) punishments ensure the sender will not deviate much from σ^* on average, and (ii) they can do no better than make an average assessment of his credibility in reporting to them. By contrast, when credibility was required at each history of the game S 's payoffs were instead no greater than under truth-telling.

It is worth noting here that while the use of a review aggregator makes the full set of commitment payoffs achievable, it does not generally provide a full folk theorem. Long-run receivers can be coerced to take dominated actions through the threat or promise of future continuation payoffs. By contrast, short-run receivers can only be persuaded via their beliefs. Hence, persuasion can never induce them to play strictly dominated actions.¹⁷

One last point we should discuss here, is the nature of the review aggregator's commitment. First, we have introduced less commitment power than is common in mechanism design. The ex-ante *uninformed* aggregator only commits to the way it garbles information that would have been public otherwise i.e. the history. By contrast, the commitment power required by the privately informed sender to his experimental design is stronger. It is as if an agent in a mechanism design problem could commit to some particular communication strategy independent of his ex-post incentives. Second, in our case, the review aggregator's behaviour could be enforced by public revelation of the past history at the end of every \mathcal{G} -phase. Receivers can then verify whether the aggregator has adhered to its rule.

Discussion: Alternative Mechanisms

One might wonder whether there are other simpler statistical tests that could be conducted and still achieve high payoffs. For example, perhaps we need not make use of receiver feedback in our tests at all? For instance, in a principal-agent setting, Jackson and Sonnenschein [2007] show that having many independent copies of an underlying mechanism design problem can allow a designer to implement the full set of incentive compatible allocation rules without transfers. In essence, they provide a budget of reports for the agent to allocate across problems, with the reports reflecting the underlying distribution of types. In our setting, this would amount to providing the sender with

¹⁷The astute reader may have noticed that our results on review aggregation do not rely on any special feature of communication games. Indeed, these results should extend to general repeated games with long and short-run players more broadly.

a budget of *Buy* messages to be used in a \mathcal{G} -phase, for example $\frac{2}{3}T$. However, a key difference in our paper is that the sender’s preferences are not naturally ordered in his private information by a single crossing condition. Indeed, given such a budget our (impatient) sender would prefer to use his entire allocation of *Buy* messages on the first $\frac{2}{3}T$ of receivers, regardless of their respective θ_i . Knowing this, a receiver who was told to *Buy* would immediately discount it as babbling, regardless of whether the aggregator garbles the history of play.

In a recent paper with long-run receivers, Margaria and Smolin [2015] describe how a folk theorem can be obtained even when the long-run receiver receives no feedback of her utility, ω_t . Their main insight is to apply the discounted quota of Frankel [2016], along with phases of rewards and punishments, to support a folk theorem by making the sender indifferent between all reports at all stages. In our setting, this is not possible without a review aggregator. This said, we believe that a review aggregator may make a similar mechanism possible. Note, however, for any such mechanism to work the aggregator would need perfect knowledge of δ . Even the slightest mistake would cause the entire mechanism to unravel. For instance, if the sender was marginally more impatient than expected, he would strictly prefer *Buy* to *Not Buy* at every instant of the mechanism. Of course, receiver IC constraints would again be violated everywhere. By contrast, our mechanism does not require such precise knowledge of δ .

4.1 Application to e-Commerce

A stylized application of the model in Section 2 helps shed light on the economic value of these review aggregators, and when they might make buyers, sellers, or both, better off. We will continue the analysis using the seller-customer model¹⁸ but where there is a unit mass of sellers, indexed by $i \in [0, 1]$, and of customers, $j \in [0, 1]$. Each customer is long-lived but can be treated as short-lived because they are anonymous and rematch with a different seller each period, never interacting with the same seller twice. An e-commerce firm, *RA*, acts as a review aggregator charging sellers a fixed share of their surplus per

¹⁸The results of this section continue to go through in a model in which the Seller can choose prices each period, so long as the threat of potential competition prevents him from capturing all the buyer’s expected surplus. For example, suppose S chooses prices in each period but there is a long-run potential entrant, whose goods are all of quality $\theta^E = 3$, faces production costs of $d^E = 2$ and an entry cost $F > 0$. Each round, S ’s costs inflate with probability δ to $d_t^S = 2$ and then remain there forever. There is an equilibrium of this game in which S signals he is low cost by maintaining a price below 2 each period, so long as his costs remain $d^S = 1$. Moreover, this yields a v function such that the seller and customer behave as above. A quantitatively different but qualitatively similar v function obtains if a risk averse customer has the outside option of buying a “brand new”, i.e. high quality for sure, product at an exogenous outside price. Alternatively, following Rhodes and Wilson [2017], we could use standard monopolists - this would have different implications for optimal information structures but would provide the same motivation for review aggregation.

transaction. Therefore RA 's profits are:

$$\pi_{RA} = \mathcal{S} \cdot \mathbb{E}[(1 - \delta) V]$$

where \mathcal{S} is the measure of sellers and \mathcal{C} the measure of customers who trade on the website. Sellers' and customers participations are determined by:

$$\mathcal{S} = (\mathbb{E}[(1 - \delta) V]^{\varepsilon_s} \cdot \mathcal{C}^\gamma)$$

$$\mathcal{C} = (\mathbb{E}[u(a, \theta)]^{\varepsilon_c} \cdot \mathcal{S}^\gamma)$$

$\varepsilon_s > 0$, $\varepsilon_c > 0$, $0 \leq \gamma < 1$ measure the responsiveness of sellers' (customers') participation to their surplus per transaction and access to customers (sellers), respectively. Notice that if $\gamma = \varepsilon_s = \varepsilon_c$, then \mathcal{S} is just an increasing, concave function of the total profits a seller expects to earn on the website. As $\gamma < 1$, we implicitly assume some decreasing returns constraining the rate of entry - which could come from offline administrative, congestion or even direct competitive costs. Finally, if we interpret sales as clicks, the firm's objective function can be motivated by the optimal pricing scheme in Baye et al. [2011].

A straightforward application of Theorem 3 tells us that the e-Commerce firm, RA , can induce an equilibrium with payoffs approximating any allocation of expected average utilities (u, v) on the Pareto frontier (PT in Figure 3). For example, a test that implements average payoffs close to truth-telling, T , is passed when bad reviews, $\omega_t = 1$, follow *Buy* messages at a rate only slightly more than $1 - p < 1/2$. Alternatively, to implement maximal persuasion average payoffs, P , the test is passed when bad reviews follow *Buy* messages at a rate slightly less than $1/2$. The profit-maximizing solution lies along PT , and satisfies

$$\frac{1 + \varepsilon_s - \gamma^2}{\gamma \varepsilon_c} = \frac{v^*}{u^*}. \quad (6)$$

Equation (6) tells us that RA 's choice of statistical test depends on two things: (i) the customers' and sellers' relative elasticities of participation with respect to their expected surplus on the platform, $\varepsilon_c/\varepsilon_s$; and (ii) the complementarity between seller and customer participation, γ . Low relative elasticity of customer to seller participation, $\varepsilon_c/\varepsilon_s$, and low complementarity between seller and customer participation, γ , tilts RA 's choice of test towards the seller's interest i.e. the allowable conditional rate of bad reviews moves towards $1/2$, shifting (u^*, v^*) towards P . When these values are high, however, the test tilts towards the customers' interests i.e. the allowable conditional rate of bad reviews moves towards $1 - p$, shifting (u^*, v^*) towards T .

This simple example provides clear comparative statics on when the use of ratings systems by e-commerce firms can improve outcomes for buyers, sellers, or both. Under bilateral trade, consider the (constrained) efficient equilibrium payoffs T' in Figure 3.




Figure 3: Payoff Sets for Online Trading

When $p \approx \frac{1}{2}$, T' is close to $(0, 0)$. In this case, the presence of an e-Commerce firm is likely to improve outcomes for everyone, regardless of $\frac{\varepsilon_R}{\varepsilon_S}$ or γ . By contrast as $p \rightarrow 1$, $T' \rightarrow T$. Accordingly, e-commerce platforms are almost certain to make sellers better off at the expense of customers. Finally, for intermediate noise, sellers always gain (weakly) from having e-commerce firms involved, whereas customers gain only if their participation is sufficiently important for RA 's profits.

4.2 Application to Financial Advice & Disclosure Rules

Many investment brokerages (the sender) offer investment advice to clients (receivers). Several papers show that brokerages have incentives to ‘oversell’ products to their clients and are far too optimistic in their recommendations [Dugar and Nathan, 1995, Lin and McNichols, 1998, Michaely and Womack, 1999, Krigman et al., 2001, Hong and Kubik, 2003].

In 2002, the National Association of Security Dealers (NASD), a financial industry self-regulating body, imposed rules that require brokerages to disclose the aggregate distribution of their recommendations to clients. Barber et al. [2006] and Kadan et al. [2009] analyze the advice given by brokerages as well as the price reaction to that advice. Prior to the introduction of these rules, analysts gave ‘Buy’ recommendations 60% of the time.¹⁹ On introduction of these new rules ‘Buy’ calls dropped almost immediately to 51% and by the following year they made up only 42% of total recommendations. Along with the drop in “Buy” recommendations Kadan et al. [2009] show that prices became *more responsive* to ‘Buy’ calls, suggesting clients found the new reports more persuasive.

A simple version of our model can shed light on this intervention. Suppose there are L clients, each of whom is considering buying an idiosyncratic (i.i.d.) portfolio that, given

¹⁹This was not anomalously high. Since the beginning of their dataset, ‘Buy’s had always exceeded 60%.

their preferences, is overvalued or undervalued: $\omega_t^l \in \{Low, High\}$. Client payoffs from buying the portfolio is $(-1/p)$ if $\omega_t^l = Low$, and $1/p$ if $\omega_t^l = High$; not-buying has payoff 0. Ex-ante, $Pr(\omega_t^l = High) = \mu_0 = 1/3$ for all t .^{20,21} The brokerage observes a signal θ_t^l on the value of each portfolio l , where $Pr(\omega_t = \theta_t | \theta_t) = p > 0.5$. The brokerage then sends a vector of messages $\boldsymbol{\mu}_t = (\mu_t^1, \dots, \mu_t^L)$ to clients. Hence, clients' expected payoffs conditional on μ_t^l are as in Section 1, and the brokerage always wants clients to buy.

Prior to the 2002 rules, client l only observed the history of their own interactions with the brokerage. Thus, this is just L replications of the standard case so the equilibrium payoff set is identical to that of Section 1: payoffs bounded strictly away from the frontier. Our predictions are consistent with the data - brokerages were attempting persuasion, over-recommending 'Buy' and were being punished with periods of babbling, in which their 'Buy' reports were having little market impact.²²

However, after 2002, it became possible for clients to observe a measure of the history of *all* messages sent by the brokerage across clients/assets. For concreteness, suppose clients observe the aggregate proportion of 'Buy' calls sent by the brokerage in the past in addition to their own advice. The proportion of 'Buys' can now act as a sort of review aggregator where the brokerage is given a per period budget of 'Buy' calls to allocate across clients. Being indifferent about which client buys, there is an equilibrium in which he recommends 'Buy' when $\theta_t^l = High$ and then allocates the remaining budget of 'Buy's at random to the remaining clients. When L is large this washes out the noise in the brokerages predictions and when the budget is less than or equal to two thirds then it is incentive compatible for the clients to follow the advice. This is all sustained by the threat of moving to a babbling equilibrium in the future.

In accordance with the data, the policy change allowed for more convincing persuasion, and less reversion to periods of babbling. In this way brokerages became more disciplined, and so more convincing, in their 'Buy' recommendations. Interestingly, from our earlier discussions on e-commerce, while the policy change likely made brokerages better off, the effect on clients' payoffs is ambiguous.

It is worth discussing some of the theoretical deviations from our review aggregator, and then the relation of these deviations to the actual problem facing brokerages and clients. First, the NASD legislation provides a finer aggregate rating than the review aggregator. As the clients are given the precise outcome of the statistical test every period the equilibrium is not selected by the review aggregator - leading to a greater

²⁰This assumption can be relaxed. It is sufficient to have ex-post observability of overall valuations across assets. While this example deserves its own asset pricing model, it is beyond the scope of the present paper.

²¹Alternatively, the sender gives advice on L asset classes, clients focus on a given class, and do not monitor the advice given on other assets. Such a model could be grounded in rational inattention or costly information acquisition. Empirically, different clients do in fact focus on different asset classes.

²²Babbling with 'Buy' signals is even more natural in a world with some naive clients who take advice at face value.

degree of multiplicity. Given this multiplicity we needn't have moved to a different equilibrium at all, we could even have moved to one worse for both senders and receivers - this is a well known problem and we don't try to solve it here. However, if the new equilibrium is outside the pre-reform equilibrium set then this would lend support to our theory regarding review aggregation.²³

Second, the NASD legislation, unlike our review aggregator, is an unconditional test - it does not make use of the realized outcomes of prices. We are able to do this in our model because the analyst is indifferent across the different portfolios. In reality, brokerages will care differently about different portfolios, moreover they make these recommendations over a period of time and so will be discounting. As noted in the discussion of alternative mechanisms above, this can cause problems for a mechanism that uses only the aggregate budget and not a conditional budget. That said, we also abstracted away from the capacity of brokerages to randomize the timing and rate of advice. Moreover, they will also have privately known heterogenous valuations across clients and portfolio. In such a setting it is still feasible to generate some uncertainty for clients about the precise strategy that they are facing and we can recover some of the effective randomization necessary to sustain a review aggregator. It is not clear without a fuller model, including macro shocks and asset pricing, how well the brokerages can do this.

Given these issues, we would tentatively suggest that brokerages may be better served by a legislation with features closer to our review aggregator. That is, i) a coarser rating where movement from bad ratings to good ratings are more insensitive to brokerage behavior; ii) the review window for each brokerage is private knowledge for the brokerage and the regulator; iii) a conditional budget for advice. Point (i) reduces the multiplicity of equilibria by enforcing stage game punishments for bad behavior. Point (ii) allows us to more effectively randomize the subjective order of clients within a review phase. Point (iii) stops the regulator being dependent on perfect knowledge of the analysts preferences across portfolios and clients. Whether this would also be beneficial for clients would depend on what rules the regulator chose for updating these ratings.

5 Optimal Information Structure

In this section, we study how long-run incentives affect the sender's optimal choice of information structure in the absence of mechanisms such as the coin and cup or review aggregator. To do this we focus on comparing the solutions to equations (Commitment) and (Long-Run Persuasion). Letting Λ^{KG} , Λ^{LP} , represent the sets of optimal information structures under commitment and long-run persuasion respectively, an immediate corollary of Theorem 1 is:

²³It is not immediately clear to us how one would run this test using the actual reform - but it could certainly be tested in the lab.

Corollary 2. *If Λ^{KG} contains a partitional information structure, then $\Lambda^{LP} \subseteq \Lambda^{KG}$.*

Corollary 2 is intuitive: if a partitional information structure is optimal under commitment, then the sender can also achieve $\hat{v}(\mu_0)$ in the long-run persuasion game (Theorem 1). Moreover, every $\lambda^{LP} \in \Lambda^{LP}$ must achieve the payoff $\hat{v}(\mu_0)$ and thus $\Lambda^{LP} \subseteq \Lambda^{KG}$. Motivated by this comparison, one might wonder whether optimality of partitional information structures is also necessary for $\lambda^{LP} \in \Lambda^{KG}$. However, in general this need not be so. Indeed, recall the leading example of Section 1. While the payoff from implementing λ^{KG} was lower without full commitment, we showed that it was nonetheless (weakly) optimal for the sender to use $\lambda^{LP} = \lambda^{KG}$.

The difference between incentives under long-run persuasion and commitment is most clear when v is differentiable. Consider Figure 4. In this case, λ^{KG} is not partitional – under commitment, it is optimal to induce a distribution over receiver posteriors μ_l^* , μ_h^* , where $\mu_h^* \in (0, 1)$. In particular, the optimal policy ensures that the chord connecting $(\mu_l^*, v(\mu_l^*))$ to $(\mu_h^*, v(\mu_h^*))$ satisfies tangency with v at μ_h^* . Given Lemma 2, this tangency condition implies that μ_h^* never lies at an interior maximum of v . However, under long-run persuasion the sender earns strictly lower payoffs from this information structure. Since both μ_l^* and μ_h^* are sent when $\theta = 0$, on-path punishments are required to ensure indifference in this state. The sender’s expected payoffs are therefore as if he were reporting μ_h^* only when $\theta = 1$ (point v' in Figure 4). But the long-run persuader can do better than this. Indeed, since he only retains the payoff from a high message in state $\theta = 1$ he should simply maximize $v(\mu_h)$ – that is, λ^{LP} induces a distribution over μ_l^* and μ_h^{**} . Indeed for this reason, whenever $N = 2$ and v is differentiable a long-run persuader’s optimal information structure always differs from the commitment optimum whenever the latter is not partitional.




Figure 4: Optimal information structures with and without commitment

We now generalize this observation. To do this we continue to abstract from the boundary issues introduced by discontinuities, as they hinder comparative statics analysis:

Assumption 2. $v(\mu)$ is differentiable, $\forall \mu \in \Delta\Theta$.

Under differentiability, a change in the sender’s marginal incentives leads directly to a change in his preferred information structure. Moreover, given uncertainty about the preferences of others, such an assumption is natural. To be concrete, consider extending the seller-customer model so that the customer enjoys a random utility from her outside option, $u \sim N(0, \sigma^2)$. The customer then buys if and only if $\mathbb{E}[\theta|\mu] - q = 2\mu - 1 \geq u$, providing the seller with expected payoffs of $v(\mu) = \Phi((2\mu - 1)\sigma^{-1})$. It is easy to see that, as $\sigma \rightarrow 0$, the problem tends to the original step function.²⁴

Of course, Corollary 2 continues to apply even when v is differentiable. To study the relationship between Λ^{LP} and Λ^{KG} when partitions are not optimal, we first study value functions v for which each $\lambda^{KG} \in \Lambda^{KG}$ has the property of being “free”:

Definition 3. A collection of posteriors $\{\mu_j\}_{j=1}^M$ are *free* if for each j , there exists an ε -ball in $\Delta\text{supp}(\mu_j)$ such that for all μ'_j in that ball, we have

$$\mu_0 \in \text{co} \left(\left\{ \{\mu_j\}_{k \neq j}, \mu'_j \right\} \right).$$

When posteriors are ‘free’ it is feasible to make local changes to any posterior independently, without violating Bayes plausibility. Freeness of λ^{KG} is an assumption on the geometry of the value function, v : recall that the posteriors $\{\mu_j^*\}_{j=1}^M$ are those used to support the concave envelope, \hat{v} , at μ_0 . In particular, we are not imposing anything on the nature of λ^{LP} . λ^{KG} will be free whenever we can use N distinct values, μ_j^* , $j = 1, \dots, N$, to support the concave envelope. Indeed, the freeness property arises in many solved examples in the literature: for instance, it arises in all binary state applications where the sender wishes to share some information, such as in Brocas and Carrillo [2007], Gill and SgROI [2012], Che et al. [2013], Gentzkow and Kamenica [2014], Perez-Richet [2014], Che and Hörner [2015], Alonso and Câmara [2016a], Rhodes and Wilson [2017], Lipnowski and Ravid [2017], and in the binary action environments of Kamenica and Gentzkow [2011], Alonso and Câmara [2016b].

By contrast, partitional information structures (other than truth-telling) are not free. Partitional structures are found in Gentzkow and Kamenica [2016a], Kolotilin [2015], Kolotilin et al., Kremer et al. [2014], where the receiver’s behavior depends only on the expectation of a continuous random variable, as well as in Ely [2017]. However, from Corollary 2, the optimal information structures are the same across the two problems in these cases.

²⁴Kolotilin et al. considers a sender who can offer menus of signals to the Receiver. These issues are orthogonal to our main results, and so we abstract from them here.

In this environment, it turns out that ‘quasi-partitional’ information structures are key to understanding the difference between the optimal choice of information structure with and without commitment:

Definition 4. An information structure λ is *quasi-partitional* if $\{\mu_j\}_{j=1}^M$ can be divided into mutually exclusive subsets $\mathcal{M}_1, \mathcal{M}_2$ with the following properties:

1. With respect to each other, the supports, $\text{supp } \mu_j, \mu_k \in \mathcal{M}_1$, form a partition of Θ .
2. For any $\mu_j \in \mathcal{M}_2, v(\mu_j) \geq v(\mu_k)$ for all $\mu_k \in \mathcal{M}_1$ such that $\text{supp}(\mu_k) \cap \text{supp}(\mu_j) \neq \emptyset$.

Notice that in long-run persuasion, there is always a dichotomy between two classes of posterior. First, there are *payoff relevant* posteriors, $\mu' \in \mathcal{M}_1$, which directly enter into the sender’s objective function. They are payoff relevant in the sense that, for each such message, there exists a state θ_i in which $\mu' \in \arg \min v(\mu)$, subject to $\mu \in M^i$. This implies that local changes in μ' directly affect the sender’s average payoffs. Second, the remaining messages, $\mu'' \in \mathcal{M}_2$, are payoff irrelevant: that is, for every $\theta_i \in \Theta, \mu' \notin \arg \min v(\mu)$, s.t. $\mu \in M^i$. For these posteriors, local movements have no direct effect on payoffs. A quasi-partitional information structure is thus one in which the *payoff relevant* messages form a partition of the state space: no two payoff relevant messages are ever sent in the same state.

Given the above, we can now provide the appropriate counterpart to Corollary 2 – a set of necessary conditions for a commitment structure $\lambda^{KG} \in \Lambda^{KG}$ to also be optimal under long-run persuasion:

Proposition 2. *Suppose that $\lambda^{KG} = (\lambda_j^*, \mu_j^*)_{j=1}^M \in \Lambda^{KG}$ induces free posteriors. Then there exists a $\lambda^{LP} \in \Lambda^{KG}$ only if there exists a quasi-partitional λ^{KG} , such that each $\mu_j^* \in \mathcal{M}_1$ maximizes v on $\mu \in \Delta \text{supp}(\mu_j^*)$.*

Our necessary conditions relate closely to partitional information structures. Indeed, Proposition 2 tells us that the only way a long-run persuader could optimally choose to replicate some λ^{KG} is if: (i) the payoff-relevant messages under λ^{KG} form a partitional information structure (with respect to each other), and (ii) these messages achieve the *global* maximum of any posterior distribution defined over the same support. Otherwise, the optimal information structures with and without commitment must be different.

The proof of Proposition 2 exploits variational arguments necessary for optimization of each problem. Recalling from Lemma 2 that $v(\mu_j^*)$ are all distinct, we argue that λ^* must be quasi-partitional in two steps: First, consider maximizing the objective function in (Long-Run Persuasion). We argue that for each payoff-relevant $\mu_j^*, v(\mu_j^*)$ is a local maximum of v on $\Delta \text{supp}(\mu_j^*)$ (the simplex over θ_i such that $\mu_j^{*,i} > 0$). Otherwise, there must exist nearby $\tilde{\mu}_j \in \Delta \text{supp}(\mu_j^*)$ such that $v(\tilde{\mu}_j) > v(\mu_j^*)$. Thus, the sender’s payoffs

increase in any $\theta_i \in \text{supp}(\mu_j^*)$ for which $v(\mu_j^*) < v(\mu_k^*)$, $\forall \mu_k$ such that $\theta^i \in \text{supp}(\mu_k)$. This strictly increases the objective in (Long-Run Persuasion). Moreover, since $\tilde{\mu}_j \in \Delta \text{supp}(\mu_j^*)$ the objective function is otherwise unchanged. Finally, freeness implies such a variation is feasible.

Now consider maximizing the objective in (Commitment). We argue that all payoff-relevant messages must have non-overlapping supports. Otherwise, there would be two payoff-relevant messages such that $\text{supp}(\mu_j^*) \cap \text{supp}(\mu_k^*) \neq \emptyset$. Suppose (without loss) that $v(\mu_j^*) > v(\mu_k^*)$. Since both messages are sent with positive probability in some state θ_i , we can always marginally increase the probability of sending message j in state θ_i at the expense of sending message k less. While this variation affects both μ_j^* and μ_k^* , they are each local maxima and therefore the first-order payoff consequences are 0. Thus, payoffs change due solely to the increase in the probability of sending message j and offsetting decrease in probability of message k . Since the sender prefers the former, this would strictly increase his payoffs. To avoid this contradiction, \mathcal{M}_1 must therefore be partitional.

Finally, we show that each $\mu_j^* \in \mathcal{M}_1$ must be actually achieve global maximum payoffs on $\Delta \text{supp}(\mu_j^*)$. This follows from $v(\mu_j^*)$ being a local maximum and an observation about tangency conditions required of the concave envelope of v .

Comparing Tradeoffs: long-run Persuasion vs. Commitment

In general, one can construct value functions v for which the concave envelope is supported by posteriors which are not free. Nonetheless, the variational arguments behind the proof of Proposition 2 can be extended to yield necessary conditions for an information structure to be optimal both with and without commitment. To illustrate, we provide heuristic derivations of the sender's first-order conditions for optimality in each case and use these to illuminate how the tradeoffs differ across problems.

Consider an interior variation $(d\lambda_j, d\mu_j)_{j=1}^M$, such that $(\lambda_j^* + d\lambda_j, \mu_j^* + d\mu_j)_{j=1}^M$ satisfies Bayes plausibility.²⁵ Then $\lambda^* \in \Lambda^{KG}$ only if the first-order condition

$$\underbrace{\sum_{j=1}^M \lambda_j^* Dv(\mu_j^*)^T d\mu_j}_{\text{Intensive Margin}} + \underbrace{\sum_{j=1}^M v(\mu_j^*) d\lambda_j^*}_{\text{Extensive Margin}} = 0 \quad (7)$$

holds, where $Dv(\mu_j) = \left(\frac{\partial v}{\partial \mu_j^1}, \dots, \frac{\partial v}{\partial \mu_j^N} \right)$ is the gradient of $v(\mu_j)$ with respect to each μ_j^i , $i = 1, \dots, N$. For a sender with commitment, equation (7) shows that a marginal change in the information structure has two consequences. First, there are changes at the

²⁵While it is easy to explicitly write the feasibility constraints, the specifics are not useful for our argument. Importantly however, so long as $(\lambda_j^*, \mu_j^*)_{j=1}^M$ is not partitional, there always exist feasible interior variations.

Extensive Margin: by altering the frequency of messages sent in state θ_i , he can enjoy the extra payoffs from the more valuable message more often. For example, suppose the sender wishes to increase the probability of sending some message μ_j in a state θ_i , by reducing the probability of sending some other message μ_k in that state, where $v(\mu_j) > v(\mu_k)$. All else equal, the sender will capture the difference $v(\mu_j) - v(\mu_k)$ with marginal probability $d\lambda_j$.

However, there is also a second effect of such a policy, at the *Intensive Margin*: By sending μ_j more (μ_k less) in state θ_i , the equilibrium values of μ_j and μ_k change, such that the sender's payoff experiences the first-order effect $\frac{\partial v}{\partial \mu_j^i} d\mu_j^i + \frac{\partial v}{\partial \mu_k^i} d\mu_k^i$ across *all* states in which it is sent with positive probability.

Similarly, we can write the necessary first order conditions for $\lambda^* \in \Lambda^{LP}$ as follows. In addition to satisfying Bayes plausibility, the variations $(d\lambda_j, d\mu_j)_{j=1}^M$ must now be chosen to ensure that we do not create discontinuous reductions in \underline{v}_i for some state θ_i (since this would be strictly sub-optimal at the margin). Indeed so long as we focus on variations that leave message supports unchanged, $\mu_j + d\mu_j \in \text{supp } \mu_j$, $j = 1, \dots, M$, we avoid costly discontinuities.²⁶ For any such variation, optimal long-run persuasion imposes the first-order constraints

$$\underbrace{\sum_{j=1}^M \tilde{\lambda}_j Dv(\mu_j^{**})^T}_{\text{Intensive Margin}} d\mu_j = 0, \quad (8)$$

where $\tilde{\lambda}_j := \sum_{i: \underline{v}_i = v(\mu_j)} \mu_0^i$ is the probability with which the sender's payoff (net of punishments) is $v(\mu_j^{**})$.^{27,28} Compared with (7), equation (8) appears much reduced. In particular, if the sender is constrained by long-run credibility concerns it turns out that he only evaluates changes at the intensive margin. This follows because of the need to ensure his long-run credibility in equilibrium. Since his payoff in each state is pinned down to \underline{v}_i , he does not experience any benefit at the extensive margin from sending higher value messages more in state θ_i . Instead, he cares only about what such deviations achieve at the intensive margin - that is, to the inherent values of the messages in \mathcal{M}_1 .

There is a second difference between (7) and (8): the weights the sender uses to evaluate the magnitude of the intensive margin. The reason for this is again the presence of on-path punishments: under long-run persuasion, the sender 'overweights' the importance of messages in \mathcal{M}_1 , since these are the only posteriors relevant to the payoffs he actually retains. By contrast, under commitment the expected value of the intensive

²⁶Again, the only information structures for which there is no feasible $(d\lambda_j, d\mu_j)_{j=1}^M$ under the additional support restriction $\mu_j + d\mu_j \in \text{supp } \mu_j$ are the partitional information structures.

²⁷Given Lemma 2, we do not have to worry about 'kinks' due to the min operator in (Long-Run Persuasion) at any $\lambda^* \in \Lambda^{KG}$.

²⁸Notice that $\tilde{\lambda}_j > 0$ if and only if $j \in \mathcal{M}_1$.

margin is determined using the actual probabilities, $\lambda_j = \Pr(m = \mu_j)$.

Since any feasible first order variation under long-run persuasion is also feasible under commitment, we can directly compare conditions (7) and (8). Thus, $\lambda^{KG} = \lambda^{LP}$ only if

$$\sum_{j=1}^M (\lambda_j^* - \tilde{\lambda}_j) Dv(\mu_j^*)^T d\mu_j + \sum_{j=1}^M v(\mu_j^*) d\lambda_j^* = 0. \quad (9)$$

Equation (9) helps to illustrate the intuition behind Proposition 2. It shows that a necessary condition for $\lambda^{KG} = \lambda^{LP}$ is that differences between the weights that a long-run and a committed sender respectively attribute to the intensive margin must exactly offset the costs the committed sender feels at the extensive margin, for *all* feasible variations. However, these intensive and extensive effects are very different from each other and there is no reason to think a priori that they would be the same. For example as the discussion at the beginning of this section noted, when $N = 2$ and v is differentiable, $\Lambda^{LP} \cap \Lambda^{KG} = \emptyset$ if and only if Λ^{KG} contains no partitional information structure.

6 Further Applications and Discussion

In our first example, we show in the context of ‘fake news’ how conclusions about the value of communication mechanisms also depends on the existing notoriety of media outlets. We then show how the use of *private* messages sent to receivers playing a public goods game can approximate the coin and cup and make everyone strictly better off than under public messaging. Our final example is the use of public record by central banks in adopting persuasion as a policy instrument.

Social Media and Fake News

There have been recent moves to tackle so called “fake news” on social media platforms.²⁹ A current proposal is to use fact-checkers such as Politifact to directly verify each story as it occurs. However, this may prove very costly to keep track of in real time. Alternatively, one could treat story providers in a similar manner to the way that eBay treats sellers and give ratings conditional on their record of veracity. While a full investigation into these issues is beyond the scope of the present paper, our simple model does highlight some important welfare considerations for those attempting to deal with ‘fake news’.

It turns out that the effects of such a policy would be very different for well known content providers, such as traditional news organizations, versus relatively unknown providers, such as someone producing stories in their mother’s basement. Well-known providers already have long-run incentives because there is an (imperfect) public record of

²⁹For instance, Allcott and Gentzkow [2017] study fake news story about presidential candidates in the 2016 election and find that about half of the people who remember a fake news story believed it.

their past veracity (contrast National Enquirer vs. New York Times). However, unknown providers have little to discipline their reporting. A reduced-form for poor observability of the record of veracity of alternative providers is to consider senders with differing discount rates, δ . By keeping track of the record of veracity for the unknown providers we can increase their effective discount rate and make it possible to sustain credibility by giving these unknown providers some credibility to gain or lose.

Remark 1. In general, any mechanism that improves the ability of the senders to persuade receivers in the future allows us to sustain meaningful communication with lower δ . This can be done by improving the monitoring of past actions or by improving the sender’s ability to use non-deterministic information structures.

By acting as a mediator sites like Facebook and Twitter could allow unknown providers to move from a state of babbling to some form of informative communication. If we apply the simplified analysis used for e-commerce above outcomes improve for both the unknown senders and their receivers, regardless of the weight social media sites put on the interests of the providers versus the public. The effects on mainstream providers and their readers is ambiguous since allocations in triangle $BD'T'$ of Figure 3 are feasible for high δ without mediation schemes. Putting aside issues of competition, the mainstream providers necessarily do better whereas their readers may fare better or worse, depending on p and the social media site’s objective function.³⁰ Interestingly, our simple model suggests that those most likely to lose out from attempts to deal with ‘fake news’ are the mainstream readers.

Team Management

We adapt the model used in Hermalin [2007]; consider an even numbered team of $L > 3$ myopic workers (receivers) and one patient benevolent boss (the sender). At period t each worker exerts costly effort $e_t^l \in \{0, 1\}$ on a new project. The total state dependent output from a project at stage t is:

$$\Pi_t = \frac{L}{2}\theta_t \sum_L e_t^l$$

where $\theta_t \in \{1, 3\}$, $Pr(\theta_t = 3) = \mu_0 = \frac{1}{3}$ and we will refer to the high payoff project as good and the low payoff as bad. All workers share equally in the total output and their effort is not measured. The stage payoff of worker l is

$$u^l = \frac{\Pi}{L} - e^l.$$

³⁰The question of competition between senders is beyond the scope of this project. In a static setting, Gentzkow and Kamenica [2016b] show there is still a role for persuasion with multiple senders. With many long-run senders, we might also expect collusion on persuasive equilibria is possible in a setting like ours.

Suppose the boss wishes to maximize the sum of workers' payoffs. As in the standard case, the boss observes the state and then makes a report to the workers who then simultaneously choose their effort level. At stage t workers observe the history of outputs $\{\Pi_0, \Pi_1, \dots, \Pi_{t-1}\}$.

Note first that it is Pareto optimal for all workers to exert effort at every stage independent of the state. However, it is also easy to see that it is only individually optimal for the each worker to exert effort if their posterior probability of the project being good is $Pr(\theta_t = 3) = \mu_t \geq 0.5$. Further, because workers are myopic, high effort cannot be enforced through punishment strategies.

If the boss is only able to send a *public* message to all workers then the game is similar to the Mayor and Firm case. There is a truth-telling equilibrium where the boss reports the true state and all workers exert effort if and only if the boss says the project is good. If the output is low, $L^2/2$, after the boss sent a message stating the project was good the workers know the boss lied and we permanently move to a babbling equilibrium. We can also support non-truthful equilibria in the same way as with the mayor and firm, but in this case both boss and worker will be worse off as the boss's incentives are aligned with the average worker.

However, if the boss is able to send a vector of *private* messages to workers we can do much better. In the good state the boss always tells workers to exert effort; in the bad state he randomly selects (equiprobably) half the workers and tells them the project is good. Conditional on being told the project is good each worker has posterior $Pr(\theta_t = 3|Good) = 0.5$ and exerts effort if and only if the boss says the project is good. On the equilibrium path we only observe $\Pi \in \left\{ \frac{L^2}{4}, \frac{3L^2}{2} \right\}$, if workers observe any other output they know the boss has deviated from the equilibrium strategy and the game moves to a babbling equilibrium forever. The ability to mix over receivers approximates the effect of the coin and cup, thus Pareto dominating public announcements.³¹

Achieving higher levels of persuasion via private messaging does not rely on this being a public good. In general, all that is necessary is that there are many simultaneous actions and some form of feedback mechanism about the history of states and message vectors.³²

Central Bankers as Persuaders

Tamura [2016] shows how a central bank who can *commit* to announcements can use persuasion to improve aggregate outcomes through the expectations channel in an economy with sticky prices. In this setting the central bank chooses a monetary policy and makes

³¹Note that by allowing the boss to send individualized messages we can achieve better payoffs via cheap talk than in Hermalin [2007].

³²An important consideration for implementing private messaging systems is the incentives for receivers to share their messages. In the public good game examined above this turns out not to be an issue as workers never wish to discourage others from working. More broadly, such institutions can be sustained in workplaces by strict informational barriers such as 'Chinese Walls'.

announcements to maximize the expected utility of the representative household. Similar to the public good problem above, agents wish to be collectively persuaded but are bound by their own individual rationality and application of Bayes rule. In the model, the central bank's optimal information strategy is deterministic (but not truth-telling). Therefore by Theorem 1, the bank can sustain the same outcomes via long-run persuasion if its information at the time of announcements can be observed ex-post. Indeed, central banks have made large strides improving both contemporaneous and ex-post transparency: the Fed, BoE, BoJ and ECB have all moved to publishing the minutes of their meetings within a month.

Central bankers' information is nonetheless still not perfectly observable. Minutes are noisy signals of policy discussions. Consequently, if it were possible to aggregate these signals over time then the central bank could do better. One might consider publishing minutes less frequently, perhaps in annual blocks. However, as in our other applications, the bank would have incentives to front-load any lies at the beginning of the review phase, unravelling the system's credibility. Given Theorem 3, one solution could be to introduce stochastic review phases (from the perspective of the public). Such a system could be implemented by the use of cryptographic technology so the length of the review is ex-post verifiable by the public.

7 Conclusions and Extensions

In this paper we have characterized the optimal payoffs and information structures under long-run persuasion. Moreover, we have been able to directly compare them to persuasion with ex-ante commitment: we found that long-run persuasion is frequently inferior and requires different information structures. But, when monitoring is perfect, the situation can be retrieved with the use of a coin and a cup mechanism. However, when monitoring is imperfect we need a review aggregator if we are to approach the payoffs and information structures available to a sender with commitment.

Both the coin and cup and review aggregator provide potential micro-foundations for the commitment assumption in the persuasion and information design literature. Each achieves the set of commitment outcomes by tampering with the observed history of the sender's behavior in an asymmetric fashion. This insight allows us to achieve payoff sets superior to those identified in earlier literatures on games with short-run players. Further work should look how these mechanisms apply in the setting of general repeated games.

Several other open questions arise from this paper. How does long-run persuasion perform when information acquisition is endogenous, or when information is partially verifiable, or when monitoring is costly? How do these results extend when there are many interacting senders or receivers? How do we optimally implement the coin and cup or review aggregator when states are not independent across time? This paper then, is

a first step on the road to having a fuller understanding of persuasion and information design when commitment has to be generated endogenously.

On top of developing the theory further, there are several applied and empirical issues which bear examining in our setting. Indeed, in Subsections 4.1, 4.2 and the Online Appendix, we illustrate key issues and trade-offs in the context of e-commerce, social media, ‘fake-news’, team management, central banking, and financial advice. Yet, to apply our results to designing policy in these areas we would need a fuller and more detailed characterization of the economic features peculiar to each environment. Moreover, the cases of e-commerce, social media and finance, seem to be ripe for empirical investigation of long-run persuasion with and without review aggregation.

References

- H. Allcott and M. Gentzkow. Social media and fake news in the 2016 election. Technical report, National Bureau of Economic Research, 2017.
- R. Alonso and O. Câmara. Bayesian persuasion with heterogeneous priors. *Journal of Economic Theory*, 165:672–706, 2016a.
- R. Alonso and O. Câmara. Persuading voters. *The American Economic Review*, 106(11):3590–3605, 2016b.
- R. J. Aumann. Mixed and behavior strategies in infinite extensive games. Technical report, Princeton University, NJ, 1961.
- R. J. Aumann and M. B. Maschler. *Repeated games with incomplete information*. MIT press, 1995.
- B. M. Barber, R. Lehavy, M. McNichols, and B. Trueman. Buys, holds, and sells: The distribution of investment banks’ stock ratings and the implications for the profitability of analysts’ recommendations. *Journal of Accounting and Economics*, 41(1-2):87–117, 2006.
- M. R. Baye, X. Gao, and J. Morgan. On the optimality of clickthrough fees in online markets. *The Economic Journal*, 121(556), 2011.
- D. Bergemann and S. Morris. Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11(2):487–522, 2016.
- J. Best and D. Quigley. Honestly dishonest: A solution to the commitment problem in bayesian persuasion, 2016. University of Oxford, Nuffield College Working Paper 2016-W12.

- J. Bizzotto, J. Rüdiger, and A. Vigier. How to persuade a long-run decision maker. 2017. Working Paper.
- I. Brocas and J. D. Carrillo. Influence through ignorance. *RAND Journal of Economics*, 38(4):931–947, 2007.
- A. Chakraborty and R. Harbaugh. Persuasion by cheap talk. *American Economic Review*, 100(5):2361–2382, 2010.
- Y.-K. Che and J. Hörner. Optimal Design for Social Learning. pages 1–49, 2015.
- Y.-K. Che, W. Dessein, and N. Kartik. Pandering to persuade. *The American Economic Review*, 103(1):47–79, 2013.
- C. Dellarocas. Reputation mechanism design in online trading environments with pure moral hazard. *Information Systems Research*, 16(2):209–230, 2005.
- U. Doraszelski and J. F. Escobar. Restricted feedback in long term relationships. *Journal of Economic Theory*, 147(1):142–161, 2012.
- A. Dugar and S. Nathan. The effect of investment banking relationships on financial analysts’ earnings forecasts and investment recommendations. *Contemporary Accounting Research*, 12(1):131–160, 1995.
- J. C. Ely. Beeps. *The American Economic Review*, 107(1):31–53, 2017.
- A. Frankel. Discounted quotas. *Journal of Economic Theory*, 166:396–444, 2016.
- D. Fudenberg and D. Levine. Efficiency and observability with long run and short run players. *Journal of Economic Theory*, 62:1030135, 1994.
- D. Fudenberg and E. Maskin. On the dispensability of public randomization in discounted repeated games. *Journal of Economic Theory*, 53(2):428–438, 1991.
- D. Fudenberg, D. M. Kreps, and E. S. Maskin. Repeated games with long-run and short-run players. *The Review of Economic Studies*, 57(4):555–573, 1990.
- M. Gentzkow and E. Kamenica. Costly persuasion. *American Economic Review*, 104(5):457–462, 2014.
- M. Gentzkow and E. Kamenica. A rothschild-stiglitz approach to Bayesian persuasion. *American Economic Review*, 106(5):597–601, 2016a.
- M. Gentzkow and E. Kamenica. Competition in persuasion. *The Review of Economic Studies*, (September):1–36, 2016b.

- M. Gentzkow and E. Kamenica. Bayesian persuasion with multiple senders and rich signal spaces. *Games and Economic Behavior*, 104:411–429, 2017.
- A. Gershkov and B. Szentes. Optimal voting schemes with costly information acquisition. *Journal of Economic Theory*, 144(1):36–68, 2009.
- D. Gill and D. SgROI. Sequential decisions with tests. *Games and Economic Behavior*, 63(2):663–678, 2008.
- D. Gill and D. SgROI. The optimal choice of pre-launch reviewer. *Journal of Economic Theory*, 147(3):1247–1260, 2012.
- J. Green and N. Stokey. Two representations of information structures and their comparisons. 1978. Institute for Mathematical Studies in the Social Sciences.
- J. C. Harsanyi. Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points. *International Journal of Game Theory*, 2(1):1–23, 1973.
- B. E. Hermalin. Leading for the long term. *Journal of Economic Behavior and Organization*, 62(1):1–19, 2007.
- H. Hong and J. D. Kubik. Analyzing the analysts: Career concerns and biased earnings forecasts. *The Journal of Finance*, 58(1):313–351, 2003.
- J. Hörner, S. Takahashi, and N. Vieille. Truthful equilibria in dynamic bayesian games. *Econometrica*, 83(5):1795–1848, 2015.
- M. O. Jackson and H. F. Sonnenschein. Overcoming incentive constraints by linking decisions. *Econometrica*, 75(1):241–257, 2007.
- B. Jullien and I.-U. Park. New, like new, or very good? reputation and credibility. *The Review of Economic Studies*, 81(4):1543–1574, 2014.
- O. Kadan, L. Madureira, R. Wang, and T. Zach. Conflicts of interest and stock recommendations: The effects of the global settlement and related regulations. *Review of Financial Studies*, 22(10):4189–4217, 2009.
- E. Kamenica and M. Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.
- A. Kolotilin. Experimental design to persuade. *Games and Economic Behavior*, 90:215–226, 2015.
- A. Kolotilin, M. Li, T. Mylovanov, and A. Zapechelnyuk. Persuasion of a privately informed receiver. *Econometrica*. Forthcoming 2017.

- I. Kremer, Y. Mansour, and M. Perry. Implementing the Wisdom of the Crowd. *Journal of Political Economy*, 122(5):988–1012, 2014.
- L. Krigman, W. H. Shaw, and K. L. Womack. Why do firms switch underwriters? *Journal of Financial Economics*, 60(2):245–284, 2001.
- H.-w. Lin and M. F. McNichols. Underwriting relationships, analysts’ earnings forecasts and investment recommendations. *Journal of Accounting and Economics*, 25(1):101–127, 1998.
- E. Lipnowski and D. Ravid. Cheap talk with transparent motives. 2017. Working Paper.
- C. Margaria and A. Smolin. Dynamic Communication with Biased Senders. Working Paper, 2015.
- R. Michaely and K. L. Womack. Conflict of interest and the credibility of underwriter analyst recommendations. *Review of financial studies*, 12(4):653–686, 1999.
- E. Perez-Richet. Interim Bayesian persuasion: First steps. *American Economic Review*, 104(5):469–474, 2014.
- E. Piermont. Distributional Uncertainty and Persuasion. Working Paper, 2016.
- R. Radner. Repeated Principal-Agent Games with Discounting. *Econometrica*, 53(5):1173–1198, 1985.
- L. Rayo and I. Segal. Optimal information disclosure. *Journal of political Economy*, 118(5):949–987, 2010.
- A. Rhodes and C. M. Wilson. False advertising. 2017. Working Paper.
- R. T. Rockafellar. *Convex analysis*. Princeton university press, 1997.
- S. Tadelis. Reputation and feedback systems in online platform markets. *Annual Review of Economics*, 8:321–340, 2016.
- W. Tamura. Optimal Monetary Policy and Transparency under Informational Frictions. *Journal of Money, Credit and Banking*, 48(6):1293–1314, 2016.
- I. A. Taneva. Information design. Working Paper, 2016.

Appendix A

We leave the proofs of Lemmas 1 and 2 and Theorem 1 to the Online Appendix B.

Proof of Proposition 1

Proof. In any equilibrium, S must be indifferent at any history (h_t, θ_t^i) between all messages $\tilde{\mu} \in \text{supp}(\sigma_t(h_t, \theta_t^i))$. Since $\underline{\mu}_i(h_t) := \{\arg \min v(\mu) : \mu \in \text{supp}(\sigma_t(h_t, \theta_t^i))\}$ is by definition in the support of $\sigma_t(h_t, \theta)$, we must have that payoffs from *any* equilibrium message at this history are

$$V_t(h_t, \theta_t) = v\left(\underline{\mu}_i(h_t)\right) + \delta \mathbb{E}[V_{t+1}((h_t, \mu_t, \theta_t), \theta_{t+1})].$$

Consider the following problem:

$$\sup_{\sigma \in \Sigma} \mathbb{E}_\theta [V_t(\underline{h}_t, \theta_0)] \quad (10)$$

s.t.

$$\begin{aligned} V_{t+\tau}(\underline{h}_{t+\tau}, \theta_{t+\tau}) &= v\left(\underline{\mu}_{t+\tau}^i\right) + \delta \mathbb{E}\left[V_{t+\tau+1}\left(\left(\underline{h}_{t+\tau}, \underline{\mu}_{t+\tau}^i, \theta_{t+\tau}\right), \theta_{t+\tau+1}\right)\right], \\ \mu_0 &\in \text{co}\left(\cup_{\theta_t \in \Theta} \text{supp}(\sigma_t(h_t, \theta_t))\right), \end{aligned}$$

$\forall \underline{h}_t$ s.t. $\tilde{\mu}_t = \underline{\mu}_t^i$ at all subsequences $\underline{h}_{t+\tau'}$, $0 \leq \tau' \leq \tau$, of \underline{h}_t at which S acts. We refer to the set of continuation payoffs that satisfy all constraints in (10), by \mathcal{V} . Notice \mathcal{V} is non-empty.³³

At $t = 0$ (where $h_0 = \emptyset$), problem (10) is a relaxed version of problem (4): it only retains constraints for histories in which S has always reported the ‘worst’ current message $\underline{\mu}_{t+\tau'}^i$, among all those available in the support of his strategy at previous histories, $h_{t+\tau'}$, $\tau' < \tau$. All other constraints from (4) are dropped. Thus, the optimal value of (10) provides an upper bound on (4).

Let $V_t^*(\underline{h}_t, \theta_t)$ be the supremum achieved in problem (10) at history \underline{h}_t . From the first constraint, we must have

$$\sup_{\sigma_t(h_t, \theta), V_{t+\tau+1}} \mathbb{E}\left[v\left(\underline{\mu}_{t+\tau}^i\right) + \delta V_{t+\tau+1}\left(\left(\underline{h}_{t+\tau}, \underline{\mu}_{t+\tau}^i, \theta_{t+\tau}\right), \theta_{t+\tau+1}\right)\right] = \mathbb{E}\left[V_{t+\tau}^*(\underline{h}_{t+\tau}, \theta_{t+\tau})\right] \quad (11)$$

where the supremum is taken over feasible lotteries $\sigma_t(h_t, \theta) \in \Lambda(\mu_0)$ and feasible payoffs from the continuation equilibrium, $V_{t+\tau+1} \in \mathcal{V}$.³⁴

For any $t + \tau$, history $(\underline{h}_{t+\tau}, \theta_{t+\tau})$ and corresponding strategies $\sigma_{t+\tau}(\underline{h}_{t+\tau}, \theta_{t+\tau})$,

³³The discounted payoff from repeated play of the static babbling equilibrium at each history, $\frac{v(\mu_0)}{1-\delta}$, is feasible.

³⁴Focusing on expected continuations (rather than values conditional on θ) ensures that we do not violate the constraint $\mu_0 \in \text{co}\left(\cup_{\theta_t \in \Theta} \text{supp}(\sigma_t(h_t, \theta_t))\right)$.

$V_{t+\tau}(\underline{h}_{t+\tau}, \theta_{t+\tau})$ is maximized by choosing the highest feasible expected continuation,

$$\mathbb{E} \left[V_{t+\tau+1}^* \left(\left(\underline{h}_{t+\tau}, \underline{\mu}_{t+\tau}, \omega_{t+\tau} \right), \theta_{t+\tau+1} \right) \right].$$

Moreover, since the continuation games at histories $\left(\underline{h}_{t+\tau}, \underline{\mu}_{t+\tau}, \omega_{t+\tau} \right)$ and $\underline{h}_{t+\tau}$ are identical, the expected continuation values must be equal:

$$\mathbb{E} \left[V_{t+\tau}^* \left(\underline{h}_{t+\tau}, \theta_{t+\tau} \right) \right] = \mathbb{E} \left[V_{t+\tau+1}^* \left(\left(\underline{h}_{t+\tau}, \underline{\mu}_{t+\tau}, \omega_{t+\tau} \right), \theta_{t+\tau+1} \right) \right]$$

Substituting into (11) yields, on rearrangement:

$$(1 - \delta) \mathbb{E}_\theta [V_0^*(\theta_0)] = \sup_{\lambda \in \Lambda(\mu_0)} \sum \mu_0^i \underline{v}_i(\lambda) \quad (12)$$

Since any equilibrium value is bounded by this supremum, the first part of our result holds.

Since $\underline{v}_i(\lambda)$ is the minimum of finitely many upper semicontinuous functions (v is upper-semicontinuous and from Lemma 1, $M \leq N < \infty$), it is upper semi-continuous. Moreover, the set $\Lambda(\mu_0)$ is clearly compact. Therefore, by the Extreme Value Theorem, the maximum exists. Let the lottery that achieves this optimum be $\lambda^* \in \Lambda(\mu_0)$, with associated support $\{\mu_1^*, \mu_2^*, \dots, \mu_{N'}^*\}$, where for convenience we index such that $v(\mu_1^*) \leq v(\mu_2^*) \leq \dots \leq v(\mu_{N'}^*)$. Notice also that compactness of $\Delta\Theta$ and the USC of v jointly imply that $v(\mu)$ is bounded above by some $\bar{v} < \infty$ for all $\mu \in \Delta\Theta$.

Let v^B be the worst expected stage payoff to the sender from any equilibrium of the stage game and σ_t^B be the sender's corresponding equilibrium strategy. To aid notation, let $\underline{v}^* := \sum_{i=1}^N \mu_0^i \underline{v}(\lambda^*)$. Now define σ_t^λ as the stage game strategy that induces the lottery λ^* . From Kamenica and Gentzkow [2011], σ_t^λ exists. Consider the following strategy σ^* , such that for all t :

$$\sigma_t^* = \begin{cases} \sigma_t^\lambda & \text{if } h_t \in \underline{H}_t \\ \sigma_t^B & \text{if } h_t \in \bar{H}_t \cup H_t^B \end{cases},$$

where H_t^B is the set of histories h_t^B such that the sender reports some $\tilde{\mu}_\tau \notin \text{supp}(\sigma_\tau^*)$ at some sub-history h_τ^B . Thus histories in $\underline{H}_t \cup \bar{H}_t$ represent on-path histories. \bar{H}_t is the set of punishment periods. A punishment period commences at t if $\tilde{\mu}_{t-1} \neq \underline{\mu}^i$ at $h_{t-1} \in \underline{H}_{t-1}$ for some $\tilde{\mu}_{t-1} \in \text{supp}(\sigma_{t-1}^\lambda)$, and lasts for $K_{i,\mu} + \mathbf{1}_{b < \beta_{i,\mu}}$ periods. Where the final period of this punishment is implemented with a public randomization device $b_t \sim U[0, 1]$.³⁵ $K_{i,\mu}$ and $\beta_{i,\mu} \in [0, 1]$ are the values of k and β that solve the following equation:

³⁵This is for simplicity only, we could instead approach the proof using deterministic punishments as in Fudenberg and Maskin [1991].

$$v(\mu) - v(\underline{\mu}_i) = \sum_{\tau=1}^k \delta^\tau (\underline{v}^* - v^B) + \beta \delta^{k+1} (\underline{v}^* - v^B). \quad (13)$$

First, note that if $\underline{v}^* = v^B$ then this payoff can trivially be sustained by repeated play of σ_t^B at every history. Thus, suppose $\underline{v}^* > v^B$. Let $\underline{\delta} < 1$ be defined by the following equation:

$$\bar{v} = \frac{\underline{\delta}}{1 - \underline{\delta}} (\underline{v}^* - v^B).$$

Since $v(\mu) - v(\underline{\mu}_i) < \bar{v} < \infty$, there always exists a k and β solving equation 13 for any $\delta \geq \underline{\delta}$. To see that σ^* is an equilibrium for $\delta > \bar{\delta}$, first note at any history $h_t \in \bar{H}_t$, all choices of $\tilde{\mu}_t \in \text{supp}(\sigma_t^B)$ yield the same continuation equilibrium and therefore σ_t^B is the sender's best response at h_t . Second, by (13), the sender is indifferent between choosing $\underline{\mu}^i$ and any $\tilde{\mu}_t \in \text{supp}(\sigma_{t-1}^\lambda)$, for all i . Finally, (13) implies that deviating from σ_t^λ for some $h_t \in \underline{H}_t$ results in an expected utility loss of

$$\sum_{\tau=1}^{\infty} \delta^\tau (\underline{v}^* - v^B) > \bar{v}.$$

□

Proof of Theorem 2

Proof. We first prove the result for $\hat{v}(\mu_0)$. Defining the state variable as $(\theta_t, c_t) \in \Theta \times [0, 1]$, it is easy to show that the optimal experiment under commitment can be written as deterministic with respect to (θ_t, c_t) . It then follows from Theorem 1 that this is achievable. Finally, for any $\underline{v}(\mu_0) \leq \nu^* \leq \hat{v}(\mu_0)$, where $\underline{v}(\mu_0)$ is the sender's worst stage game payoff, we can achieve ν^* with a public randomization device, which randomizes between play of the optimal experiment and the cheap talk equilibrium. For δ large enough, this continues to be an equilibrium supported by the threat of permanently moving to the worst stage game equilibrium. □

Proof of Theorem 3

Proof. For simplicity, we establish the result for the seller-customer environment of Section 1. The Online Appendix contains a restatement of the result and its proof for the more general environment.

The proof is constructive. For any feasible information structure λ , we find a rule for transitioning between \mathcal{G} and \mathcal{B} such that, for a patient enough seller, there is an equilibrium of the review system in which: (i) on-path costs of falling into a \mathcal{B} -phase are negligible; (ii) the seller's optimal reporting strategy approximately induces λ on average within a \mathcal{G} -phase; (iii) each customer finds it optimal to follow the seller's

recommendations. To aid exposition, we show the result for $\lambda = \lambda^{KG}$. The argument is identical for other choices of λ .

Let $(q_\theta)_{\theta \in \{1,3\}} \in [0, 1]^2$ be a pair of parameters satisfying $q_1 = 0.5 - \epsilon$, $q_3 = 1$, where $\epsilon > 0$, and consider the following class of review systems: At the T^{th} period of any arbitrary \mathcal{G} -phase, \mathcal{G}_j , the review system continues on to another T period \mathcal{G} -phase (the seller ‘passes the test’) if

$$\frac{1}{T} \sum_{t \in \mathcal{G}_j} \mathbf{1}(m_t = \text{Buy}, \omega_t = \omega) \in [\mu_0 z_3^\omega q_3 + (1 - \mu_0) z_1^\omega q_1 - \chi, \mu_0 z_3^\omega q_3 + (1 - \mu_0) z_1^\omega q_1 + \chi]$$

for $\omega = 1, 3$, where $\mathbf{1}(E)$ is the indicator for event E , $z_\theta^\omega = \Pr(\omega_t = \omega \mid \theta_t = \theta)$, and $\chi > 0$ is a parameter measuring the test’s strictness.³⁶ Otherwise, the seller ‘fails’ and the review system switches to a \mathcal{B} -phase for βT periods. At the end of any \mathcal{B} -phase, the review system reverts to a new \mathcal{G} -phase.

Suppose that during any \mathcal{G} -phase customers obey the seller’s recommendation, while in any \mathcal{B} -phase customers treat reports as babbling. We first show that there exist δ^* , T^* , χ^* , β^* , such that, for all $\delta \geq \delta^*$, the seller’s best response in the review system with parameters T^* , χ^* , β^* , involves adopting a strategy which (i) passes each test with high probability; (ii) he earns discounted average payoffs close to $\frac{2}{3}(1 - \epsilon)$.

For any $p \neq 0.5$,³⁷ it follows from Chebyshev’s weak Law of Large Numbers that for any $\bar{l} > 0$, $\epsilon_2 > 0$, we can find a sequence $(\chi_n, T_n)_{n=1}^\infty \rightarrow (0, \infty)$ that (i) if the seller reports according to strategy $(q_\theta)_{\theta \in \{1,3\}}$ throughout a \mathcal{G} -phase, then $\Pr(\text{Pass}) \geq 1 - \epsilon_2$; while (ii) if

$$\frac{\sum_{t \in \mathcal{G}_j} \mathbf{1}(m_t = \text{Buy}, \theta_t = \theta)}{T} \notin [q_\theta - \bar{l}, q_\theta + \bar{l}],$$

for some $\theta \in \{1, 3\}$ then $\Pr(\text{Fail}) \geq 1 - \epsilon_2$. Fix parameters $0 < \epsilon_2^* < \bar{\varphi}$, where $\epsilon_2^*/\bar{\varphi}$ can be chosen arbitrarily small. Given any χ^* , T^* satisfying (i) and (ii) for $\epsilon_2 = \epsilon_2^*$, the seller’s payoff from any recursive strategy at the start of a \mathcal{G} -phase is³⁸

$$\mathcal{V}_{\mathcal{G}} = \sum_{t=0}^{T^*-1} \delta^t \Pr(m_t = \text{Buy}) + \delta^{T^*} (1 - \varphi + \varphi \delta^{\beta T^*}) \mathcal{V}_{\mathcal{G}} \quad (14)$$

where φ denotes the probability that the seller fails review phase \mathcal{G} . We now find β^* , δ^* such that the seller will optimally choose $\varphi < \bar{\varphi}$, if $\delta \geq \delta^*$. If the seller reports according to $(q_\theta)_{\theta \in \{1,3\}}$, his payoffs within a \mathcal{G} -phase would be $\frac{2}{3}(1 - \epsilon) \frac{1 - \delta^{T^*+1}}{1 - \delta}$, and from (i), $\varphi \leq \epsilon_2^*$. Moreover, the expected cost to the seller of failing is no greater than $\frac{2}{3}\epsilon_2^* \beta^* T^*$. Now, set $\beta^* = 4/\bar{\varphi}$ and consider any other strategy with $\varphi \geq \bar{\varphi}$. Clearly an upper bound

³⁶We suppress indices denoting specific phases for notational convenience. See Appendix B.

³⁷Assumption 3 in the Online Appendix is needed to generalize this argument. It allows statistical test to identify ‘lies’.

³⁸In the proof we establish that S has an optimal strategy, and it is in recursive strategies.

on the *incremental* payoff such a strategy can yield within the \mathcal{G} -phase is $\frac{1}{3}T^*$, while the marginal failure cost is at least

$$\delta^{T^*+1} \frac{2}{3} (1 - \epsilon) \varphi \left(1 - \frac{\varepsilon_2^*}{\bar{\varphi}} \right) \beta^* T^*.$$

For δ^* large (and $\varepsilon_2^*/\bar{\varphi}$ small) enough, it is easy to see that the expected marginal costs of failing exceed the benefits whenever $\varphi \geq \bar{\varphi}$.

Finally, notice that the expected payoffs from adopting the reporting strategy $(q_\theta)_{\theta \in \{1,3\}}$ are very close to the commitment benchmark as we take ϵ , ε_2^* small enough: from such a strategy, the seller earns a discounted average payoff of at least $(1 - \delta_n) \mathcal{V}_G \geq 2/3 - O((\varepsilon_2^*)^{-1})$. Thus, for any ϵ , $\bar{\varphi}$, and δ^* sufficiently large, we have established the existence of the appropriate T^* , χ^* , β^* .

We now show that, for a choice of $\bar{\varphi}$ sufficiently small, it is a best response for any customer to follow the seller's recommendation in a \mathcal{G} -phase. To do this, we argue that ε_2 , $\bar{\varphi}$ can be chosen small enough to ensure $|\mu_B - \mu'_B| < \varepsilon_1/2$, where

$$\mu_B := \frac{\sum_{t \in \mathcal{G}} \Pr(m_t = Buy, \theta_t = 3, t | \mathcal{G})}{\sum_{t \in \mathcal{G}} \Pr(m_t = Buy, t | \mathcal{G})} \quad (15)$$

is the equilibrium posterior of a buyer, given rating \mathcal{G} and recommendation *Buy*, and $\mu'_B = \frac{\mu_0 q_3}{2/3(1-\epsilon)} := 1/2 + \varepsilon_1$ is the belief that would be induced if the seller adopted reporting strategy $(q_\theta)_{\theta \in \{1,3\}}$. Define the seller's private history at period T of some \mathcal{G} -phase by $g_t^s = (m_\tau, \theta_\tau)_{\tau=0}^t$, the set of all such histories be G^t , and let the seller's (mixed) reporting strategy within such a phase be described by a collection of functions, $r_t(\theta_t, g_{t-1}^s) := \Pr(m_t = Buy | \theta_t, g_{t-1}^s)$, $\forall g_t^s \in G^t$, $t = 1, \dots, T$.³⁹ Given (ii) above and $p \neq \frac{1}{2}$, Chebyshev's law of large numbers implies that the seller can only fail with a probability less than some $\bar{\varphi} \rightarrow 0$ if he adopts a reporting strategy on which

$$\left| \frac{\sum_{t \in \mathcal{G}} r_t(\theta_t, g_{t-1}^s)}{T} - q_\theta \right| \leq \bar{l}$$

along histories, g_T^s , which occur with arbitrarily high probability. In the Online Appendix, we establish formally that there exist ε_2^* , $\bar{\varphi}$, $\bar{l} \rightarrow 0$ such that the seller's optimal strategy must involve $\Pr(|\sum_{t \in \mathcal{G}} r_t(\theta_t, g_{t-1}^s)/T - q_\theta| \leq \bar{l}) \rightarrow 1$ in any \mathcal{G} -phase.⁴⁰

Now consider a customer's beliefs given rating \mathcal{G} and recommendation *Buy*. We can

³⁹Notice that we have truncated histories to only include information within \mathcal{G} -phase. Given customer and review aggregator behavior, it is without loss for the purposes of our results to consider seller-optimal behavior in this way.

⁴⁰The derivation is notationally heavy without adding insight, and so relegated to the Online Appendix.

calculate the receiver's posterior probabilities in (15) as

$$\begin{aligned}
\sum_{t \in \mathcal{G}} \Pr(m_t = \text{Buy}, \theta_t = 3, t \mid \mathcal{G}) &= \sum_{t \in \mathcal{G}} \Pr(m_t = \text{Buy}, \theta_t = 3 \mid \mathcal{G}, t) \Pr(t \mid \mathcal{G}) \\
&= \sum_{t \in \mathcal{G}} \Pr(m_t = \text{Buy} \mid \mathcal{G}, t, \theta_t = 3) \Pr(\theta_t = 3 \mid \mathcal{G}, t) \Pr(t \mid \mathcal{G}) \\
&= \mathbb{E} \left[\sum_{t \in \mathcal{G}} r_t(\theta_t, g_{t-1}^s) \frac{\mu_0}{T} \right],
\end{aligned}$$

where the last line follows from (a) the aggregator's uniform permutation of customers within a \mathcal{G} -phase; (b) the i.i.d. nature of θ_t ; and (c) the definition of the seller's strategy. But since $\frac{\sum_{t \in \mathcal{G}} r_t(\theta_t, g_{t-1}^s)}{T}$ is a bounded random variable which converges to q_θ with probability approaching 1, we must also have $\mathbb{E} \left[\sum_{t \in \mathcal{G}} r_t(\theta_t, g_{t-1}^s) / T \right] \rightarrow q_\theta$. Thus,

$$\sum_{t \in \mathcal{G}} \Pr(m_t = \text{Buy}, \theta_t = 3, t \mid \mathcal{G}) \rightarrow \frac{1}{3}$$

as $\bar{\varphi} \rightarrow 0$ and similar calculations imply

$$\sum_{t \in \mathcal{G}} \Pr(m_t = \text{Buy}, t \mid \mathcal{G}) \rightarrow \frac{2}{3} (1 - \epsilon).$$

Thus, for $\bar{\varphi}$ small, $\mu_B \rightarrow \frac{1}{2(1-\epsilon)} > \frac{1}{2}$. Thus, customers indeed obey a $m = \text{Buy}$. It is easy to verify that they also obey $m = R$.

Finally, babbling can obviously be sustained as an equilibrium during \mathcal{B} -phases, since it is an equilibrium of the stage game.

Since we are free to choose $\epsilon > 0$ as small as we necessary, we have therefore established that payoffs and the underlying information structure from λ^{KG} can be approximated arbitrarily closely. \square

Proof of Proposition 2

Proof. Let $\mathcal{M} = \{\mu_j^*\}_{j=1}^M$ and for each state θ_i , $M^i := \{\mu : \mu \in \mathcal{M}, \mu^i > 0\}$ be the set of messages sent in that state. Suppose that there is an information structure $\lambda^{KG} = (\lambda_j^*, \mu_j^*)_{j=1}^M$ which induces free posteriors and is optimal under both commitment and long-run persuasion. Then it must solve the following two problems:

$$\max_{(\lambda_j, \mu_j)_{j=1}^M} \sum_{\mathcal{M}} \lambda_j v(\mu_j) \quad s.t. \quad \mu_0 = \sum_{j=1}^M \lambda_j \mu_j \tag{16}$$

and

$$\max_{(\lambda_j, \mu_j)_{j=1}^M} \sum_{\mathcal{M}} \mu_0^i v_i(\lambda) \quad s.t. \quad \mu_0 = \sum_{j=1}^M \lambda_j \mu_j. \quad (17)$$

Note that for any feasible policy there trivially exists a subset $\tilde{\mathcal{M}}_1 \subseteq \{\mu_j^*\}_{j=1}^M$ such that each $\mu_k^* \in \tilde{\mathcal{M}}_1$ satisfies $\mu_k^* \in \arg \min \{v : \mu \in M^i\}$ for some $i \in \{1, \dots, N\}$ and $\cup_{\mu_k^* \in \tilde{\mathcal{M}}_1} \text{supp}(\mu_k^*) = \Theta$. We first show that since λ^{KG} maximizes (17) and $\{\mu_j^*\}_{j=1}^M$ are free, there is a subset $\hat{\mathcal{M}}_1 \subseteq \tilde{\mathcal{M}}_1$ for which (i) $\cup_{\mu_k^* \in \hat{\mathcal{M}}_1} \text{supp}(\mu_k^*) = \Theta$ and (ii) the directional derivatives satisfy

$$Dv(\mu_k^*)^T d\mu_k = 0 \quad \forall \mu_k^* \in \hat{\mathcal{M}}_1, \quad (18)$$

$\forall d\mu_k$ such that $\mu_k^* + d\mu_k \in \Delta \text{supp} \{\mu_k^*\}$. We argue by contradiction. Suppose such a $\hat{\mathcal{M}}_1$ cannot be found. Then since v is differentiable there must exist $\theta_i \in \Theta$ and, $\forall \mu_k^* \in \tilde{\mathcal{M}}_1 \cap M^i$, some feasible $\mu_k^* + d\mu_k \in \Delta \text{supp} \{\mu_k^*\}$ for which $Dv(\mu_k^*)^T d\mu_k > 0$.⁴¹ Let this set of posteriors be \mathcal{D} and write $\mathcal{M}' = \{\mu'_j\}_{j \in \mathcal{D}} \cup \{\mu_j^*\}_{j \in \{1, \dots, M\} / \mathcal{D}}$. Since by definition of M , $\lambda_j^* > 0, \forall j \in \{1, \dots, M\}$, and $\mu'_k = \mu_k^* + d\mu_k \in \text{aff} \{\mu_1^*, \dots, \mu_M^*\}$ for all feasible $d\mu_k$, there exist scalars $\{\varepsilon_j\}_{j=1}^M$ such that $\mu_0 \in \text{co}(\mathcal{M}')$, $\forall \mu'_k = \mu_k^* + d\mu_k \in \text{aff} \{\mu_1^*, \dots, \mu_M^*\}$ satisfying $|d\mu_k| \leq \varepsilon_k$. But then standard arguments imply there must exist $\{\mu'_j\}_{j \in \mathcal{D}}$ such that $v(\mu'_j) > v(\mu_j^*), \forall j \in \mathcal{D}$ and therefore

$$\sum_{l=1}^N \mu_0^l \min \{v(\mu) : \mu \in \mathcal{M}', \mu^l > 0\} > \sum_{l=1}^N \mu_0^l \min \{v(\mu) : \mu \in \mathcal{M}\}, \quad (19)$$

where $\min \{v(\mu) : \mu \in \mathcal{M}', \mu^l > 0\} \geq \min \{v(\mu) : \mu \in \mathcal{M}\}$ holds for all $l \in \{1, \dots, N\}$ and with strict inequality for $l = i$. But (19) is a contradiction to λ^{KG} as a solution to (17). Thus, such a $\hat{\mathcal{M}}_1$ exists.

We now argue that each $\mu_k^* \in \hat{\mathcal{M}}_1$ must additionally satisfy $v(\mu_k^*) > v(\mu), \forall \mu \in \Delta \text{supp} \{\mu_k^*\}$. Suppose not. Then there exists $\mu' \in \Delta \text{supp} \{\mu_k^*\}$ such that $v(\mu') > v(\mu_k^*)$. For a vector $z \in \mathbb{R}^N$, we can define the function $h(z) = v(\mu_k^*) + \frac{\|z - \mu_k^*\|}{\|\mu' - \mu_k^*\|} (v(\mu') - v(\mu_k^*))$. Notice that for any $z' = \mu_k^* + \alpha(\mu' - \mu_k^*), \alpha \in \mathbb{R}$, $(z', h(z'))$ maps the chord connecting $(\mu_k^*, v(\mu_k^*))$ to $(\mu', v(\mu'))$. Thus, the derivative of h for a marginal change $dz \propto \frac{z' - \mu_k^*}{\|z' - \mu_k^*\|}$ is

$$Dh(\mu_k^*)^T dz = \frac{v(\mu') - v(\mu_k^*)}{\|\mu' - \mu_k^*\|} > 0.$$

However, by (18), $Dv(\mu_k^*)^T dz = 0$. Thus, by standard arguments there exists $\alpha' < 0$ sufficiently small for which

$$h(\mu_k^* + \tilde{\alpha}(\mu' - \mu_k^*)) < v(\mu_k^* + \tilde{\alpha}(\mu' - \mu_k^*))$$

⁴¹This is immediate since $Dv(\mu_k^*)^T d\mu_k \neq 0$, and $\mu_k^* - d\mu_k \in \Delta \text{supp} \{\mu_k^*\}$ if $\mu_k^* + d\mu_k \in \Delta \text{supp} \{\mu_k^*\}$.

$\forall \tilde{\alpha} \in [\alpha', 0)$. Choose some such α and let $\tilde{z} = \mu_k^* + \tilde{\alpha}(\mu' - \mu_k^*)$. Thus, we can write $\mu_k^* = \left(\frac{1}{1-\tilde{\alpha}}\right)\tilde{z} + \left(-\frac{\tilde{\alpha}}{1-\tilde{\alpha}}\right)\mu'$. Since h is linear, it must therefore satisfy

$$\left(\frac{1}{1-\tilde{\alpha}}\right)h(\tilde{z}) + \left(\frac{-\tilde{\alpha}}{1-\tilde{\alpha}}\right)h(\mu') = h(\mu_k^*) = v(\mu_k^*).$$

However, since $v(\tilde{z}) > h(\tilde{z})$ and $h(\mu') = v(\mu')$, we must also have

$$\left(\frac{1}{1-\tilde{\alpha}}\right)v(\tilde{z}) + \left(\frac{-\tilde{\alpha}}{1-\tilde{\alpha}}\right)v(\mu') > v(\mu_k^*).$$

Finally, consider the following two-stage lottery over beliefs. In the first stage, generate μ_j according to λ^{KG} . Conditional on the posterior μ_k^* , instead provide the Receiver with a lottery over posteriors \tilde{z}, μ' with probabilities $\frac{1}{1-\tilde{\alpha}}, \frac{-\tilde{\alpha}}{1-\tilde{\alpha}}$, respectively. Since $\tilde{\alpha} < 0$, such a policy satisfies Bayes plausibility and non-negativity. Moreover, the sender's payoffs from such a policy are

$$\sum_{j \neq k} \lambda_j^* v(\mu_j^*) + \lambda_k^* \left(\left(\frac{1}{1-\tilde{\alpha}}\right)v(\tilde{z}) + \left(\frac{-\tilde{\alpha}}{1-\tilde{\alpha}}\right)v(\mu') \right) > \sum_{j=1}^N \lambda_j^* v(\mu_j^*)$$

– a contradiction to λ^{KG} as the solution to (16).

We now argue that, for any two posteriors $\mu_j^*, \mu_k^* \in \hat{\mathcal{M}}_1$ such that $\text{supp}(\mu_j^*) \cap \text{supp}(\mu_k^*) \neq \emptyset$, it must be the case that $v(\mu_j^*) = v(\mu_k^*)$. Suppose not. Then there exist $\mu_j^*, \mu_k^* \in \hat{\mathcal{M}}_1$ with $\text{supp}(\mu_j^*) \cap \text{supp}(\mu_k^*) \neq \emptyset$ and for which $v(\mu_j^*) > v(\mu_k^*)$. Consider the following deviation from λ^{KG} : in some state $\theta_i \in \text{supp}(\mu_j^*) \cap \text{supp}(\mu_k^*)$, increase $\lambda_j^i := \Pr(m = \mu_j | \theta_i)$ to $\tilde{\lambda}_j^i = \lambda_j^i + d\lambda$, $d\lambda > 0$, correspondingly decrease λ_k^i to $\tilde{\lambda}_k^i = \lambda_k^i - d\lambda$ and adjust the associated posteriors respectively as follows:

$$\begin{aligned} \tilde{\mu}_j^* &= \frac{\lambda_j^i}{\lambda_j^i + d\lambda} \mu_j^* + \frac{d\lambda}{\lambda_j^i + d\lambda} \mathbf{1}_i \\ \tilde{\mu}_k^* &= \frac{\lambda_k^i}{\lambda_k^i - d\lambda} \mu_k^* - \frac{d\lambda}{\lambda_k^i - d\lambda} \mathbf{1}_i \end{aligned}$$

The policy $\tilde{\lambda} = \left(\left((\lambda_l)_{l \neq j, k}, \tilde{\lambda}_j, \tilde{\lambda}_k \right), \left((\mu_l)_{l \neq j, k}, \tilde{\mu}_j^*, \tilde{\mu}_k^* \right) \right)$ is Bayes Plausible, and satisfies non-negativity so long as $d\lambda < \lambda_k^i$. Moreover, taking derivatives in $d\lambda$ yields the necessary FOC for λ^{KG} to maximize the objective function (16)

$$\lambda_j^* Dv(\mu_j^*)^T \frac{d\tilde{\mu}_j^*}{d\lambda} - \lambda_k^* Dv(\mu_k^*)^T \frac{d\tilde{\mu}_k^*}{d\lambda} + v(\mu_j^*) - v(\mu_k^*) = 0$$

But since λ^{KG} maximizes (17), it satisfies (18) at both μ_j^* and μ_k^* , for all feasible $d\mu$. Thus, $Dv(\mu_k^*)^T \frac{d\tilde{\mu}_j^*}{d\lambda} = Dv(\mu_k^*)^T \frac{d\tilde{\mu}_k^*}{d\lambda} = 0$. Therefore, in order to maximize (16), we must also have $v(\mu_j^*) = v(\mu_k^*)$ – a contradiction to $v(\mu_j^*) > v(\mu_k^*)$.

Thus, we have established that messages in $\hat{\mathcal{M}}_1$ can be partitioned into subsets $\{s_c\}_{c=1}^{M'}$, $s_c \subseteq \hat{\mathcal{M}}_1$, with the following properties: (i) $\mu_k^*, \mu_j^* \in s_c \implies v(\mu_j^*) = v(\mu_k^*)$; (ii) $\mu_k^* \in s_c, \mu_j^* \notin s_c \implies \text{supp}(\mu_j^*) \cap \text{supp}(\mu_k^*) = \emptyset$.

By Lemma 2, the concave envelope of v is generically not supported by posteriors with equal payoffs. Thus, each s_c is generically a singleton – providing us with the required set \mathcal{M}_1 .

Finally, the properties of \mathcal{M}_2 follow trivially from the definition of $\mathcal{M}_2 = \mathcal{M}/\mathcal{M}_1$ and the fact that $\cup_{\mathcal{M}_1} \text{supp}(\mu_j^*) = \Theta$. \square

Appendix B

Below we provide the formal statement underlying Assumption 1 in the main text:

Assumption. For any $\lambda = (\lambda_j, \mu_j)_{j=1}^M \in \Delta\Theta$ such that $v(\mu_j) = v(\mu_k)$, for some $\mu_j \neq \mu_k$, $j, k \in \{1, 2, \dots, M\}$, $M > 1$, the concavification of v on $\Delta\Theta$, $\hat{v}(\mu)$, satisfies

$$\hat{v}\left(\sum_{i=1}^M \lambda_i \mu_i\right) > \sum_{i=1}^M \lambda_i v(\mu_i).$$

Proof of Lemma 1

Proof. The first part of this proof shows that it is without loss to study game $\hat{\Gamma}_\infty$ and is based on a relabelling of messages. The second part shows that it is without loss for the message space to contain at most N posteriors at each history, one fewer than is necessary under commitment (Kamenica and Gentzkow [2011]).

Take any equilibrium of game Γ^∞ , and denote the corresponding strategies and beliefs respectively by $s^* = (s_t^*(\phi_t, \theta_t))_{t=0}^\infty$, $\rho_t^*(\phi_t, m_t)$, $t = 0, 1, 2, \dots$, for S , R_t , and $\mu_t^*(\phi_t, m_t)$, $t = 0, 1, 2, \dots$. Since these strategies and beliefs form an equilibrium of Γ^∞ , they obey conditions 1. - 3. in Section 2. We construct strategies and beliefs within game $\hat{\Gamma}^\infty$ which (i) induce the same conditional distributions of receiver's actions, given appropriately defined histories, and (ii) form a direct equilibrium of $\hat{\Gamma}^\infty$; that is, they satisfy conditions 1. - 3. of the direct equilibrium definition.

We assume for ease of exposition that there is a message $\underline{m} \in M$ which is never played on the equilibrium path.⁴² For each public history $\phi_t = (m_\tau, a_\tau, \theta_\tau)_{\tau=0}^{t-1}$, define the belief induced by message m at history ϕ_t as $\tilde{\mu}_t(\phi_t, m_t) := \mu_t^*(\phi_t, m_t)$ and a corresponding on

⁴²Lemma 1 does not require such an assumption (proof available on request), but it significantly reduces the notation required to describe strategies fully.

path history of the game as $\hat{\Gamma}^\infty$, $h_t^\phi := (\tilde{\mu}_t(\phi_t, m_t), a_t, \theta_t)$. Let the set of all such on path histories be H^ϕ , where any history $h_t \notin H^\phi$ is an off-path history. Let ϕ'_t be any off path history of Γ^∞ in which $m_\tau = \underline{m}$ for any $\tau < t$.⁴³ Consider now strategies $\hat{\sigma}$, $\hat{\rho}_t$, $t = 0, 1, \dots$, and beliefs $\mu_t(h_t, \tilde{\mu}_t)$ in game $\hat{\Gamma}^\infty$, where:

$$\hat{\sigma}_t(h_t, \theta_t) := \begin{cases} \lambda_t^*(\phi_t, \theta_t), & \text{if } h_t \in H^\phi \\ \lambda_t^*(\phi'_t, \theta_t), & \text{otherwise.} \end{cases}$$

where $\lambda_t^* \in \Delta(\Delta\Theta)$ satisfies $\Pr(\hat{\mu}_t = \tilde{\mu}_t(\phi_t, m_t) \mid h_t, \theta_t) = s_t^*(m_\tau \mid \phi_t, \theta_t)$ (for $\hat{\mu}_t$ drawn from $\hat{\sigma}_t$), and $s_t^*(m_\tau \mid \phi_t, \theta_t) := \Pr(m = m_\tau \mid \phi_t, \theta_t; s_t^*(\phi_t, \theta_t))$ is the measure over M induced by lottery $s_t^*(\phi_t, \theta_t)$;

$$\hat{\rho}_t(h_t, \hat{\mu}_t) = \begin{cases} \rho^*(\phi_t, m_t), & \text{if } h_t \in H^\phi, \hat{\mu}_t = \tilde{\mu}_t(\phi_t, m_t), \text{ for some } m_t \in M \\ \rho^*(\phi'_t, m_t), & \text{if } h_t \notin H^\phi, \hat{\mu}_t = \tilde{\mu}_t(\phi'_t, m_t), \text{ for some } m_t \in M \\ \underline{a}, & \text{otherwise.} \end{cases}$$

and

$$\mu_t(h_t, \hat{\mu}_t) = \begin{cases} \mu_t(\phi_t, m_t), & \text{if } h_t \in H^\phi, \hat{\mu}_t = \tilde{\mu}_t(\phi_t, m_t), \text{ for some } m_t \in M \\ \mu_t(\phi'_t, m_t), & \text{if } h_t \notin H^\phi, \hat{\mu}_t = \tilde{\mu}_t(\phi'_t, m_t), \text{ for some } m_t \in M \\ \underline{\mu}, & \text{otherwise.} \end{cases}$$

where $a_t = \underline{a} \in A$ is a rationalizable action that minimizes S 's stage payoff and $\underline{\mu}$ is the corresponding belief.

At any history $h_t^\phi \in H^\infty$, strategy $\hat{\sigma}$ assigns probability $s_t^*(m_t \mid \phi_t, \theta_t)$ to message $\hat{\mu}_t = \mu_t(\phi_t, m_t)$. Moreover, given message $\hat{\mu}_t$, $\hat{\rho}_t(h_t^\phi, m_t)$ induces the same lottery over $a \in A$ that $\rho^*(\phi_t, m_t)$ does, conditional on (ϕ_t, m_t) . Thus, strategy profile $(\hat{\sigma}, \hat{\rho})$ induces the same distribution over receiver actions at history h_t^ϕ in game $\hat{\Gamma}^\infty$ as does (s^*, ρ^*) in Γ^∞ . Moreover, the transition probabilities between h_t^ϕ and $h_{t+1}^\phi = (h_t^\phi, \tilde{\mu}_t(\phi_t, m_t), a_t, \theta_t)$ are clearly identical to those between ϕ_t and $\phi_{t+1} = (\phi_t, m_t, a_t, \theta_t)$. Thus, $(\hat{\sigma}, \hat{\rho})$ induces the same ex ante distribution over a_t as (s^*, ρ^*) as well.

Given S 's strategy, a message $\hat{\mu}_t = \tilde{\mu}_t(\phi_t, m_t)$ induces R_t 's beliefs at any history

⁴³More formally, we need to introduce a public randomization device to $\hat{\Gamma}^\infty$. This is because it could be the case that there are two (or more) messages that induce the same belief but imply different continuation equilibria. By reducing them down to a single belief based message we would need a public randomization device to select between the two continuations.

$h_t^\phi \in H^\phi$ to be

$$\begin{aligned} \Pr(\theta_t = \theta_t^i \mid \hat{\mu}_t) &= \frac{\Pr(\hat{\mu}_t \mid h_t^\phi, \theta_t^i) \mu_0^i}{\sum_j \Pr(\hat{\mu}_t \mid h_t^\phi, \theta_t^j) \mu_0^j} \\ &= \frac{\Pr(m_t \mid \phi_t, \theta_t^i) \mu_0^i}{\sum_j \Pr(m_t \mid \phi_t, \theta_t^j) \mu_0^j} \\ &= \tilde{\mu}_t^i(\phi_\tau, m_\tau), \end{aligned}$$

$i = 1, \dots, N$. Thus, condition 2. of direct equilibrium is satisfied for such strategies. Moreover, since $\tilde{\mu}_t(\phi_\tau, m_\tau)$ are posterior probabilities, they naturally integrate back to the prior as in Kamenica and Gentzkow [2011]. Thus, condition 3. is satisfied.

We close this part of the proof by arguing condition 1. of a direct equilibrium also holds. Trivially, $\hat{\rho}_t(h_t, \hat{\mu}_t)$ is a best response for R_t to signal $\hat{\mu}_t = \tilde{\mu}_t(\phi_t, m_t)$ sent by S , since $\rho^*(\phi_t, m_t)$ was optimal for him in game Γ^∞ given these same beliefs. Moreover, $\hat{\sigma}$ is preferred for S at any history h_t than any strategy $\hat{\sigma}'$ in which $\hat{\mu}_t = \tilde{\mu}_t(\phi_t', m_t)$, for some $t, m_t \in M$, since both such strategies are relabelings of s^* and some alternative feasible strategy s' in game Γ , respectively (and similarly for $\hat{\rho}, \rho^*$). Since s^* is optimal in game Γ^∞ , $\hat{\sigma}$ is preferred to $\hat{\sigma}'$. Finally for any other deviation, we have specified the strategies for the sender and receiver in such a way that the stage payoff from the sender deviation could only be lower than that of a deviation to some alternative strategy $\hat{\sigma}'$ in which $\hat{\mu}_t = \tilde{\mu}_t(\phi_\tau', m_\tau)$, for some $\tau, m_\tau \in M$, and moreover, the continuation payoff would be the same as under $\hat{\sigma}'$.

Next, we show that we can restrict attention to message space $M \leq N$. Suppose for a contradiction that for some equilibrium payoff $\mathbb{E}[V(h_\tau, \theta_\tau)]$ of the sender and some history h_τ , the minimum number of messages in the sender's strategy compatible with obtaining $\mathbb{E}[V(h_\tau, \theta_\tau)]$ in equilibrium is $|M'| = N' > N$, where $M' = \cup_{\theta \in \Theta_t} \text{supp}(\sigma_\tau(h_\tau, \theta))$. This strategy induces a N' -point distribution $\nu \in \Delta(\Delta\Theta)$ of posterior beliefs $\{\mu_\tau(m)\}_{m \in M'}$ over θ_τ and a corresponding distribution over receiver R_τ 's actions, $a_\tau(\mu_\tau(m))$, where

$$a_\tau(\mu_\tau(m)) \in \arg \max_{a \in A} \mathbb{E}[u_R(a, \theta) \mid \mu_t] = \sum_{i=1}^N \mu_\tau^i \cdot u_R(a, \theta^i)$$

For this to be an equilibrium, it must be that for all $m_\tau \in \text{supp}(\sigma_\tau(h_\tau, \theta))$ and any $\tilde{m} \in M$,

$$\begin{aligned} V(h_\tau, \theta_\tau) : &= v(\mu_\tau(m_\tau)) + \delta \mathbb{E}[V((h_{\tau+1}, \theta_{\tau+1}))] \\ &\geq v(\mu_\tau(\tilde{m})) + \delta \mathbb{E}\left[V\left(\left(\tilde{h}_{\tau+1}, \theta_{\tau+1}\right)\right)\right] \end{aligned}$$

where $h_{\tau+1} = (h_\tau, m_\tau, a_\tau, \theta_\tau)$ and $\tilde{h}_{\tau+1} = (h_\tau, \tilde{m}, \tilde{a}_\tau, \theta_\tau)$. In particular, given any state θ_τ

and messages $m_\tau, \tilde{m}_\tau \in \text{supp}(\sigma_\tau(h_\tau, \theta))$, we must have

$$v(\mu_\tau(m_\tau)) + \delta \mathbb{E}[V((h_{\tau+1}, \theta_{\tau+1}))] = v(\mu_\tau(\tilde{m}_\tau)) + \delta \mathbb{E}\left[V\left(\left(\tilde{h}_{\tau+1}, \theta_{\tau+1}\right)\right)\right]$$

Given any history, we define an equilibrium message $m^\theta \in M'$ to be *uniquely prescribed* at state θ if $\text{supp}(\sigma_\tau(h_\tau, \theta)) = \{m^\theta\}$. The set of all messages that are uniquely prescribed at some state $\theta \in \Theta_\tau$ is denoted M^θ . We divide the set of equilibrium messages sent at history h_τ into two mutually exclusive and exhaustive sub-groups: those that are uniquely prescribed, $m \in M^\theta$, and those that are not, $m \in M'/M^\theta$.

Since $N' > N$, there exists an $\tilde{m} \in M'$ and corresponding $\mu_\tau(\tilde{m}) \in \{\mu_\tau(m)\}_{m \in M'}$ that can be removed from the support such that remaining posteriors still satisfy Bayes' plausibility

$$\sum_{m_\tau \in M'/\{\tilde{m}\}} \alpha_{m_\tau} \mu_\tau(m_\tau) = \mu_0 \quad (20)$$

for some weights α_{m_τ} such that $\alpha_{m_\tau} \geq 0$, $\sum \alpha_{m_\tau} = 1$ (follows from Caratheodory's Theorem applied to the convex set, $\Delta(\Delta\Theta)$). By Proposition 1 in Kamenica and Gentzkow [2011], posteriors $\mu_\tau(\tilde{m}) \in \{\mu_\tau(m)\}_{m \in M'/\{\tilde{m}\}}$ can be sustained by a feasible signal structure with $N' - 1$ distinct messages. Moreover, the message \tilde{m} cannot be uniquely prescribed in any state $\theta \in \Theta$. Otherwise, there would exist some θ^i for which $\mu_\tau^i(m) = 0$, $\forall m \in M'/\{\tilde{m}\}$, while $\mu_0^i > 0$, violating (20). Therefore, $\tilde{m} \in M'/M^\theta$ and for every state θ in which σ proscribes $\Pr(m_\tau = \tilde{m} | h_\tau, \theta) > 0$, there exists another message m'_θ sent with positive probability in state θ .

Construct a new strategy σ^* which induces the distribution $(\alpha_{m_\tau})_{m_\tau \in M'/\{\tilde{m}\}}$ over the posteriors $\{\mu_\tau(m)\}_{m \in M'}$ at history h_τ , and plays according to σ otherwise (this is feasible, by Proposition 1 of Kamenica and Gentzkow [2011]). For any $m \in M'/\{\tilde{m}\}$, the strategy continues to induce belief $\mu_\tau(m)$ at history h_τ and leaves continuation payoffs unchanged at $V(h_\tau, \theta_\tau)$ thereafter (for any $\theta_\tau \in \Theta_\tau$). Moreover, this continuation payoff is well defined for each m since \tilde{m} was never uniquely prescribed.

Therefore, strategy σ^* achieves the same payoffs for the sender from history h_τ (due to indifference across all messages at that history), leaves payoffs otherwise unchanged at other histories, and involves only $N' - 1$ messages sent at history h_τ . Therefore, it also does not affect incentive compatibility of equilibrium play at any prior history, h_t , for $t < \tau$. It trivially does not affect the incentive compatibility of any history h_t , for $t > \tau$. But this is a contradiction to N' as the minimum number of messages in any strategy consistent with $\mathbb{E}[V(h_\tau, \theta_\tau)]$. Finally, we note that an almost identical argument to the above establishes it is without loss to restrict attention to at most $N + 1$ -point distributions at each h^t , for characterizing the set of attainable equilibrium payoff profiles for the sender and *all* receivers. \square

Proof of Lemma 2

Proof. Fix μ_0 . Let U be the set of all real-valued upper-semicontinuous functions on $\Delta\Theta$ such that $v(\mu_0) < \hat{v}(\mu_0)$, with typical member $v \in U$, and consider the metric space $(U, \|\cdot\|)$ endowed with the sup norm. As in the text, we denote the concavification of v by \hat{v} , and an element of $\Delta\Theta$ by μ . We show that the set U^* , defined as

$$U^*(\mu_0) = \left\{ v \in U : \hat{v}(\mu_0) > \sum_j \lambda_j v(\mu_j), \forall \lambda \in \Lambda(\mu_0) \text{ s.t.} \right. \\ \left. v(\mu_j) = v(\mu_k), \text{ for some } \mu_j \neq \mu_k, j, k = 1, \dots, N+1 \right\}$$

is open and dense.⁴⁴

To establish density of U^* , consider a function $v' \in U/U^*$. We show that there exist arbitrarily small perturbations of v' under the sup norm such that the perturbed function, v , lives in U^* .

First, we establish a useful properties of the concavification of v' , which can be expressed

$$\hat{v}' = \sup_{\lambda \in \tilde{\Lambda}(\mu_0)} \sum_{i=1}^N \lambda_i v'(\mu_i)$$

where $\tilde{\Lambda}(\mu_0)$ is the subset of N -point distributions in $\Lambda(\mu_0)$.⁴⁵ Notice that any $\lambda \in \tilde{\Lambda}(\mu_0)$ can be represented as a point in $\mathbb{R}^{(N)^2}$, since we can write $\lambda = (\lambda_j, \mu_j)_{j=1}^{N+1}$. Moreover, $\tilde{\Lambda}(\mu_0)$ is by definition a closed convex polytope, is bounded by the hypercube $[0, 1]^{(N)^2}$, and is therefore a compact subset of $\mathbb{R}^{(N)^2}$. Since v' is upper semi-continuous, the function $\sum_{i=1}^N \lambda_i v'(\mu_i)$ therefore attains its maximum on $\tilde{\Lambda}(\mu_0)$. In what follows, we make use of the below Lemma:

Lemma 3. $\hat{v}(\mu_0)$ is continuous in $\mu_0 \in \Delta\Theta$.

Proof. Suppose not. Then there exists $\mu'_0 \in \Delta\Theta$ and a sequence $\{\mu_0^n\}_n$, with $\mu_0^n \rightarrow \mu'_0$, such that $\hat{v}(\mu'_0) \neq \lim \hat{v}(\mu_0^n)$. Suppose first that $\hat{v}(\mu'_0) > \lim_{\mu_0^n \rightarrow \mu'_0} \hat{v}(\mu_0^n)$. Consider some arbitrary $\tilde{\mu}_0$ and consider evaluating \hat{v} along the line segment $\mu_0^\alpha = \alpha\mu'_0 + (1-\alpha)\tilde{\mu}_0$. From the properties of concave envelopes it follows that $\hat{v}(\mu_0^\alpha) \geq \alpha\hat{v}(\mu'_0) + (1-\alpha)\hat{v}(\tilde{\mu}_0)$ for any $\alpha \in [0, 1]$. Taking limits as $\alpha \rightarrow 1$, we must have $\lim \hat{v}(\mu_0^\alpha) \geq \hat{v}(\mu'_0)$. Since this is true for any $\tilde{\mu}_0 \in \Delta\Theta$, there cannot exist a sequence $\{\mu_0^n\}_n$, with $\mu_0^n \rightarrow \mu'_0$, such that $\hat{v}(\mu'_0) > \lim \hat{v}(\mu_0^n)$ – a contradiction.

To rule out $\hat{v}(\mu'_0) < \lim_{\mu_0^n \rightarrow \mu'_0} \hat{v}(\mu_0^n)$, notice that $\Lambda(\mu_0)$ is a closed convex polytope in (μ_0, λ) and therefore compact. Since v is an upper semi-continuous function, a straightforward extension of the extreme value theorem to upper-semicontinuous functions implies that $\hat{v}(\mu'_0) \geq \lim_{\mu_0^n \rightarrow \mu'_0} \hat{v}(\mu_0^n)$. \square

⁴⁴By Caratheodory's Theorem, it is without loss of generality to define U^* with regard to finite sets, $\{\mu_i\}_{i=1}^{N+1}$.

⁴⁵Again, Caratheodory's Theorem implies this is without loss.

Thus, the subgraph of \hat{v}' ,

$$\text{sub}(\hat{v}') = \{(\mu, \nu) : \nu \leq \hat{v}'(\mu), \mu \in \Delta\Theta\},$$

is a closed (Lemma 3) convex (Kamenica and Gentzkow [2011]) set. Bound $\text{sub}(\hat{v}')$ below by some $\underline{B} \in \mathbb{R}$, such that $\underline{B} < \min_{\mu \in \Delta\Theta} \hat{v}'(\mu)$ and define the bounded, closed convex set, $H(\hat{v}') := \text{sub}(\hat{v}') \cap \{(\mu, \nu) : \nu \geq \underline{B}\}$. Note that $H(\hat{v}' + \epsilon)$ has the same properties and $H(\hat{v}') \in \text{int}(H(\hat{v}' + 2\epsilon))$.

We are now able to find an ϵ -perturbation of v' such that the new function v satisfies $v \in U^*$. Partition $\Delta\Theta$ into two sets: $C = \{\mu : v'(\mu) = \hat{v}'(\mu)\}$ and $\bar{C} := \Delta\Theta/C$. We now construct a polyhedral convex set $P := \text{co}\left(\{(\mu_j^*, \nu_j^*)\}_{j=1}^J\right)$ for which $\text{int}(H(\hat{v}')) \subset P \subset \text{int}(H(\hat{v}' + \epsilon))$ and the subset of vertices of P for which $\nu_i^* > \hat{v}'(\mu)$,

$$Q := \{(\mu_j^*, \nu_j^*) : \nu_i^* > \hat{v}'(\mu)\},$$

satisfy $\mu_j^* \in C$.

To achieve the required perturbation, note first that for any $x \in H(\hat{v}')$, we can choose a simplex S_x such that $x \in S_x$ and $S_x \in \text{int}(H(\hat{v}' + 2\epsilon))$. Because $\sum_{i=1}^{N+1} \lambda_i v(\mu_i)$ attains its maximum on $\Delta\Theta$, for each $x \notin \text{int}(H(\hat{v}'))$ ⁴⁶ we can in fact choose S_x such that at least one of its vertices $(s_1, s_2, \dots, s_{N+1})$ satisfies $s_i = (\mu'_i, \hat{v}'(\mu_i) + \epsilon)$, where $\mu'_i \in C$ for all s_i . From the union of these simplexes, $\bigcup \{S_x\}$, we can find a finite subset of simplexes whose convex hull also covers $H(\hat{v}')$ - this is the polyhedron P (Theorem 20.4, Rockafellar, 1997). Moreover, by construction, no vertex of P that lies everywhere above $H(\hat{v}')$ has $\mu_i^* \in \bar{C}$. Therefore, since $\text{int}(H(\hat{v}')) \subset P$ and P is convex, $\max_{\nu} \{\nu : (\mu, \nu) \in P\}$ must be attained by a convex combination of vertices in P taken exclusively over the subset Q .

Define the perturbed function \tilde{v} by adding ϵ to \hat{v}' at each vertex $\{(\mu_i^*, \nu_i^*)\}_{i=1}^M$ of P . \tilde{v} is clearly still upper semi-continuous. Moreover, the concavification of \tilde{v} is indeed $\max_{\nu} \{\nu : (\mu, \nu) \in P\}$. On \tilde{v} , it suffices to check that $\nu_i^* \neq \nu_j^*$ for $i \neq j \in \{1, \dots, M\}$. If two such i, j can be found, we can find a perturbation of ν_i^* by some $\tilde{\epsilon}$ satisfying $0 < \tilde{\epsilon} < \epsilon$, such that $\nu_i^* \neq \nu_k^*$, $k \in \{1, 2, \dots, M\} / \{i\}$ and the new polyhedron P' still contains C everywhere. This new perturbed function v is upper semi-continuous, satisfies $v \in U^*$ and $|v - \hat{v}'| < (M + 1)\epsilon$, which can be chosen arbitrarily close to 0.

We now show that U^* is open in v, μ_0 . Specifically, we show that for any function $v \in U^*$, there exist $\epsilon_1, \epsilon_2 > 0$ s.t. for all \tilde{v} satisfying $\|\tilde{v} - v\| < \epsilon_2$, $\tilde{\mu}_0$ satisfying $|\tilde{\mu}_0 - \mu_0| < \epsilon_1$, we have $\tilde{v} \in \bigcup_{|\tilde{\mu}_0 - \mu_0| < \epsilon_2} U^*(\mu_0)$.

Take some $v \in U^*$. Before considering perturbed functions, we first show that, for

⁴⁶For any $x \in \text{int}(H(\hat{v}'))$, surrounding simplexes such that $s_i \in \text{int}(H(\hat{v}'))$ can be trivially found. However, these vertices will never be part of the upper envelope of set P and we thus do not discuss them in detail.

some $\delta_1 > 0$, there exists $\tilde{\epsilon}_1, \tilde{\epsilon}_2 > 0$ such that $\forall |\tilde{\mu}_0 - \mu_0| < \tilde{\epsilon}_1$, if $|\hat{v}(\tilde{\mu}_0) - \sum_i \lambda_i v(\mu_i)| < \delta_1$ for some $\lambda \in \Lambda(\tilde{\mu}_0)$ then $|v(\mu_i) - v(\mu_j)| \geq \tilde{\epsilon}_2$ for all $i \neq j, i, j \in \{1, 2, \dots, N+1\}$.⁴⁷ To derive a contradiction, suppose this were not the case. Then, for any $\delta, \tilde{\epsilon}_1, \tilde{\epsilon}_2 > 0$, we could find some $\tilde{\mu}_0$ and $\lambda \in \Lambda(\tilde{\mu}_0)$ such that (i) $|\tilde{\mu}_0 - \mu_0| < \tilde{\epsilon}_1$, (ii) $|v(\mu_j) - v(\mu_k)| < \tilde{\epsilon}_2$, for some $j \neq k$, (iii) $|\hat{v}(\tilde{\mu}_0) - \sum_j \lambda_j v(\mu_j)| < \delta_1$. Now consider any sequence $(\delta^n, \tilde{\epsilon}_1^n, \tilde{\epsilon}_2^n)_{n=1}^\infty$ satisfying $\lim_{n \rightarrow \infty} (\delta^n, \tilde{\epsilon}_1^n, \tilde{\epsilon}_2^n) = 0$. Thus, we can find a corresponding sequence $((\tilde{\mu}_0^n, \lambda_0^n)_{n=1}^\infty$ in which each $(\tilde{\mu}_0^n, \lambda_0^n)$ satisfies (i)-(iii) evaluated at $\delta = \delta^n$, $\tilde{\epsilon}_1 = \tilde{\epsilon}_1^n$ and $\tilde{\epsilon}_2 = \tilde{\epsilon}_2^n$. But since $(\mu_0, \lambda) \in \mathbb{R}^N \times \mathbb{R}^{N+1}$ and $\Lambda(\mu_0)$ is compact in (μ_0, λ) , the Bolzano-Weierstrass Theorem implies that we can find a convergent subsequence $((\tilde{\mu}_0^{n'}, \lambda_0^{n'})) \rightarrow (\mu_0, \lambda^*)$ for some $\lambda^* \in \Lambda(\mu_0)$. Moreover, upper semi-continuity of v implies that at this limit, we must either have (i) $\hat{v}(\mu_0) = \sum_i \lambda_i^* v(\mu_i^*)$, and $v(\mu_j^*) = v(\mu_k^*)$, for some $j, k \in \{1, 2, \dots, N+1\}$; (ii) $\sum_i \lambda_i^* v(\mu_i^*) > \hat{v}(\mu_0)$, or (iii) $\hat{v}(\mu_0) = \sum_i \lambda_i^* v(\mu_i^*)$, $v(\mu_i^*) \neq v(\mu_j^*)$, $\forall i, j \in \{1, 2, \dots, N+1\}$.⁴⁸ Since $v \in U^*$, case (i) yields a contradiction. By definition of \hat{v} , case (ii) also implies a contradiction. Finally, we rule out case (iii). In this case there must be a discrete upward jump in v at μ_j^* , for some j , must also cause a discontinuity at $\hat{v}(\mu_0)$ on the path $\mu_0^n \rightarrow \mu_0$ - a contradiction to the continuity of \hat{v} , which we proved above.

Finally, consider any perturbed function v' such that $\|v' - v\| \leq \min\{\frac{\delta}{3}, \frac{\tilde{\epsilon}_2}{3}\}$. Using properties (i) - (iii) of v above and standard triangle inequality arguments it is straightforward to find $\tilde{\epsilon}'_2, \delta'_1 > 0$ such that, if we have

$$\left| \hat{v}'(\tilde{\mu}_0) - \sum_i \alpha_i v'(\mu_i) \right| < \delta'_1,$$

for some $\lambda \in \Lambda(\tilde{\mu}_0)$ s.t. $|\tilde{\mu}_0 - \mu_0| < \tilde{\epsilon}_1$, then $|v'(\mu_i) - v'(\mu_j)| \geq \tilde{\epsilon}'_2$ for all $i \neq j, i, j \in \{1, 2, \dots, N+1\}$. Thus, U^* is indeed open. \square

Proof of Theorem 1

Proof. (If) Clearly, the optimal discounted average payoff achievable via information design on each receiver R_t weakly exceeds the optimal payoff from any repeated game (since this problem is similar to problem (4)), but without incentive constraints). Suppose that at prior μ_0 , the optimal payoff under information design, $\hat{v}(\mu_0)$, can be implemented by a bijection \hat{s}_P between some partition P of Θ to $M := \{m_1, m_2, \dots, m_M\}$, where $M \leq N$. Thus, for each $\theta^i \in \Theta$, $\hat{s}(\theta^i) = m(\theta^i)$, for some unique $m \in M$. Moreover, we can define an inverse function $m^{-1}(m_j) := \{\theta : m(\theta) = m_j\} \subset \Theta$, with the property that $m^{-1}(m_j) \cap m^{-1}(m_k), \forall j, k \in \{1, 2, \dots, M\}, j \neq k$ and $\cup_{j \in \{1, \dots, M\}} m^{-1}(m_j) = \Theta$. Under

⁴⁷Again, by Caratheodory's Theorem it is without loss to restrict attention to $N+1$ -point distributions, $\lambda \in \tilde{\Lambda}(\tilde{\mu}_0)$.

⁴⁸Since \hat{v} is continuous, we cannot have at the limit $\hat{v}(\mu_0) > \sum_i \alpha_i v(\mu_i)$.

such a strategy, a receiver's posterior belief, conditional on observing a message $m_j \in M$ is a vector $\mu(m_j)$, where the i^{th} entry of μ is

$$\mu^i(m_j) = \Pr(\theta \mid \theta \in m^{-1}(\Theta))$$

S 's payoff from experiment (M, \hat{s}_P) is

$$\sum_{i \in \{1, \dots, N\}} \mu_0^i v(\mu(m_j(\theta_i)))$$

Now, consider the repeated cheap talk game and the following lottery, λ_P , whose support is $\{\mu(m_j)\}_{j \in \{1, 2, \dots, M\}}$. Under lottery λ_P ,

$$\Pr(\mu = \mu(m_j)) = \sum_{\theta^i \in m^{-1}(m_j)} \mu_0^i$$

Lottery λ_P replicates the induced distribution of posteriors under \hat{s}_P : therefore, it is clearly feasible, $\lambda_P \in \Lambda(\mu_0)$. Moreover, since each θ^i induces one and only one message under λ_P , $\underline{v}_i(\lambda_P) = v(\mu(m_j(\theta)))$. Therefore,

$$\sum \mu_0^i \underline{v}_i(\lambda_P) = \sum_{i \in \{1, \dots, N\}} \mu_0^i v(\mu(m_j(\theta)))$$

Since the optimal payoff from information design is an upper bound on that under repeated persuasion, λ_P must achieve the maximum value of (5).

Finally, by Proposition 1 there exists a $\underline{\delta} < 1$ such that we can obtain this payoff as an equilibrium of the repeated game for all $\underline{\delta} \leq \delta < 1$ - establishing necessity.

(*Only if*) Suppose that $\hat{v}(\mu_0)$ cannot be achieved by a partitional information structure. Take any optimal *non-partitional* experiment (M, s^{**}) and denote the corresponding information structure $\lambda^{**} = (\lambda_j^{**}, \mu_j^{**})_{j=1}^M$. The expected payoff from this experiment is

$$\hat{v}(\mu_0) = \sum_{j \in M} \lambda_j^{**} v(\mu_j^{**}) = \sum_{i \in \Theta} \mu_0^i \left(\sum_{j \in M} \frac{\lambda_j^{**} \mu_j^{i, **}}{\mu_0^i} v(\mu_j^{**}) \right)$$

However, by definition of $\underline{v}_i(\lambda^{**})$ we have

$$\underline{v}_i(\lambda^{**}) \leq \sum_j \frac{\lambda_j^{**} \mu_j^{i, **}}{\mu_0^i} v(\mu_j^{**}) \tag{21}$$

We now argue that inequality (21) is strict for some $i \in \{1, 2, \dots, N\}$, and hence average payoffs from the information structure under long-run persuasion are lower than under

commitment:

$$\sum_{i \in \Theta} \mu_0^i v_i(\lambda^{**}) < \hat{v}(\mu_0). \quad (22)$$

To see this, suppose not. Then by the definition of “non-partitional” there exists some state i in which two or more messages are sent. Consider two messages m and m' sent in the same state, it follows that equation (21) holds with equality only if these two messages have the same stage payoff: $v(\mu(m)) = v(\mu(m'))$. This contradicts Assumption 1.⁴⁹ Hence, inequality (21) is strict and inequality (22) then follows.

Finally, note that inequality (21) holds for any information structure λ' , not only those induced by optimal experiments. Hence, as $\hat{v}(\mu_0)$ is only achieved by $\lambda \in \Lambda^{KG}$ and all λ^{KG} are non-partitional by assumption then

$$\sum \mu_0^i v_i(\lambda') < \hat{v}(\mu_0)$$

for all λ' . □

Appendix C

General Theorem on Mediated Communication Under Imperfect Monitoring

To establish the main result of this section, we make a common assumption on the form of relationship between ω_t and θ_t . Denote the joint probability of (θ_t^i, ω_t^j) , $i, j \in \{1, 2, \dots, N\}$ by $f(\theta^i, \omega^j) := \mu_0^i g(\omega_t^j | \theta_t^i)$. Then:

Assumption 3. *There exists a subset $\Omega' := \{\omega^1, \dots, \omega^N\} \subset \Omega$ such that*

$$(f(\theta_t^i, \omega_t^1))_{i=1}^N, \dots, (f(\theta_t^i, \omega_t^N))_{i=1}^N$$

are linearly independent.

Assumption 3 requires that while the receiver cannot typically infer the sender’s signal θ_t from observing her own ω_t , each possible signal provides statistically distinct posterior distributions $\theta | \omega^i$ in the sense that such a posterior could be formed only by ω^i and not some linear combination of signals, $\{\omega^j\}_{j \neq i}$.

In addition, we need the following regularity assumption on the sender’s reduced-form preferences:

Assumption 4. *For all $\lambda' \in \Lambda(\mu_0)$, there exists an open set $X(\lambda') \subset \Lambda(\mu_0)$ such that $\sum_{i=1}^N \lambda_i v(\mu_i)$ is continuous in λ on $X(\lambda')$ and $\lambda' \in \text{cl}(X(\lambda'))$.*

⁴⁹Recall, this assumption follows from the non-genericity result in Lemma 2.

Assumption 4 ensures that the sender's reduced-form utility function does not involve discontinuous jumps in the sender's payoffs which are only attainable on a lower-dimensional subset of beliefs in $\Delta\Theta$. Since u_S is continuous, such 'jumps' only occur if the receiver's beliefs just induce him to change his actions in a discontinuous way, on a 'knife edge' region in the space of posterior beliefs. We rule out this kind of knife-edge preference of the receiver to take discontinuously high actions for two related reasons. First, such utility functions are clearly not robust to small perturbations of receiver's preferences. Second, they are not robust to small amounts of noise in receiver's posterior beliefs. Finally, to avoid some technical issues that do not add any insight to the main result, we assume in the rest of this section that A is a countable set.

Before introducing the Theorem we define a Review Mechanism as follows

Definition of Review Mechanisms

These mechanisms are implemented by a *Mediator*, who commits ex ante to play a pre-specified role as we describe below.

Let $\mathcal{G}(j)$, $\mathcal{B}(j)$, $j = 1, 2, 3, \dots$ be defined inductively by:

$$\mathcal{B}(j) \in \{0, \beta T\} \quad (23)$$

$$\mathcal{G}(j+1) = \mathcal{G}(j) + T + \mathcal{B}(j) \quad (24)$$

for each $j \in \mathbb{N}_{>0}$, where $\beta \in \mathbb{R}_{>0}$, $T \in \mathbb{N}_{>0}$, are parameters such that $\beta T \in \mathbb{N}_{>0}$, and

$$\mathcal{G}(1) = 0 \quad (25)$$

We define the (T -period) *review phase* \mathcal{G}^j as the set of time periods

$$\mathcal{G}^j = \{\mathcal{G}(j), \mathcal{G}(j) + 1, \dots, \mathcal{G}(j) + T - 1\}.$$

If $\mathcal{B}(j) > 0$, then we say that review phase j is followed by an associated *punishment phase* $\mathcal{B}^j := \{\mathcal{G}(j) + T, \mathcal{G}(j) + T + 1, \dots, \mathcal{G}(j) + T + \mathcal{B}(j)\}$; otherwise ($\mathcal{B}(j) = 0$), review phase j is punishment-free and we write $\mathcal{B}^j := \{\emptyset\}$. Let the set of histories of phases $(\mathcal{G}_j, \mathcal{B}_j)_{j=1}^t$ satisfying (23)-(25), $t = 1, 2, \dots$, and $t \in \mathcal{G}_t \cup \mathcal{B}_t$ be \mathcal{A}_t . Finally, let $\pi_j : \mathcal{G}^j \rightarrow \mathcal{G}^j$ be an arbitrary permutation function (bijection) on \mathcal{G}^j and denote the set of all such permutation functions by $\Pi(\mathcal{G}^j)$.

The Mediator is able to design constraints on the history of play that receivers can observe. In particular, at the beginning of any review phase \mathcal{G}^j , the Mediator randomly selects $\tilde{\pi} \in \Pi(\mathcal{G}^j)$ from the uniform distribution, where $\Pr(\pi_j = \tilde{\pi}) = \frac{1}{|\Pi(\mathcal{G}^j)|}$. Given $\tilde{\pi}$, the Mediator commits to permute the ordering of receivers $R_{\mathcal{G}(j)}, \dots, R_{\mathcal{G}(j)+T-1}$ by assigning receiver R_τ to play the mechanism at period $\tilde{\pi}(\tau)$, for $\tau \in \mathcal{G}^j$. At time τ , he

privately informs receiver $R_{\tilde{\pi}^{-1}(\tau)}$ of *only* the current phase of the mechanism, \mathcal{G}^j . Receivers do not learn the realization of $\tilde{\pi}$ and are unable to see their position in the line, $\tilde{\pi}(\tau)$, directly.⁵⁰ Similarly, the Mediator uniformly permutes the ordering of receivers in punishment phases.

At each time $t \in \{0, 1, \dots\}$, the review mechanism asks the sender to make a *private* report, $\tilde{\theta}_t \in \Theta$, to the Mediator. At each time t at which he is asked to report, the sender observes the history of his own signals, $(\theta_\tau)^t = (\theta_1, \dots, \theta_t)$, reports, $(\tilde{\theta}_\tau)^{t-1} = (\tilde{\theta}_1, \dots, \tilde{\theta}_{t-1})$, and the history of phases announced by the Mediator, $(\mathcal{G}_j, \mathcal{B}_j)_{j=1}^t$ for some ι such that $t \in \mathcal{G}_\iota \cup \mathcal{B}_\iota$. Given a report of $\tilde{\theta}_t$ in period t , the Mediator then privately sends a message $m \in M$ to receiver $R_{\tilde{\pi}^{-1}(t)}$ according to a (possibly random) message function $r : \mathcal{G}^\iota \cup \mathcal{B}^\iota \times \Theta \rightarrow \Delta M$, where we note in particular the message function depends on whether the mechanism is in a review phase or a punishment phase at period t . Finally, receiver $R_{\tilde{\pi}^{-1}(t)}$ observes message m and chooses an action a from the set A .

To be clear, in a *review mechanism* the sender's strategy is a collection of (possibly mixed) reporting functions $\tilde{\theta}_t : \Theta^t \times \Theta^{t-1} \times \mathcal{A}_{t-1} \times \mathcal{P}_t \rightarrow \Delta\Theta$, for $t = 0, 1, 2, \dots$ where $\mathcal{P}_t := \cup_{\iota: t \in \mathcal{G}_\iota \cup \mathcal{B}_\iota} \mathcal{G}^\iota \cup \mathcal{B}^\iota$. In other words, at each time, the sender can condition his reports on the entire history of play he has observed in the mechanism, including his own past types, reports and the Mediator's announcements of past and present phases. A strategy for receiver R_t is a function $\rho : \mathcal{P}_t \times M \rightarrow A$, in which she chooses an action as a function of the currently announced phase and the message m sent to her by the Mediator.

To close the description of our review mechanisms, we need to define the circumstances under which review phase \mathcal{G}_j is followed by a punishment, $\mathcal{B}(j) > 0$, for $j = 1, 2, \dots$. To do this, consider some review phase \mathcal{G}_j . Given any history of receiver signals $(\omega_\tau)^{\mathcal{G}(j)+T-1} = (\omega_1, \omega_2, \dots, \omega_{\mathcal{G}(j)+T-1})$, reports $(\tilde{\theta}_\tau)^{\mathcal{G}(j)+T-1}$ and any $\omega^k \in \Omega'$, define the subsequence of time periods in which ω^k was realized in review phase \mathcal{G}_j as

$$\kappa(\omega^k, (\omega_\tau)^t, \mathcal{G}^j) := \{\tau : \tau \in \mathcal{G}^j \text{ s.t. } \omega_\tau = \omega^k\},$$

and the corresponding empirical frequency of reports, $\tilde{\theta}^i \in \Theta$, as:

$$\mathcal{F}\left(\tilde{\theta}^i \mid (\tilde{\theta})^{\mathcal{G}(j)+T-1}, \omega^k, \mathcal{G}^j\right) = \frac{\sum_{\tau \in \kappa(\omega^k, \mathcal{G}^j)} \mathbf{1}(\tau : \tilde{\theta}_\tau = \tilde{\theta}^i)}{|\kappa(\omega^k, \mathcal{G}^j)|}$$

With some notational abuse, we will often suppress the dependence of \mathcal{F} on $(\tilde{\theta})^{\mathcal{G}(j)+T-1}$,

⁵⁰It is convenient but not crucial for our main results that receivers observe *nothing* about the mechanism outside their own phase. It is crucial, however, that receivers cannot infer anything about their position within the *current phase* from any other information they may have. For real recommendation platforms such as eBay, uncertainty about how/when the platform updates its reviews and/or the frequency of customer interaction help to keep buyers uninformed in this way.

\mathcal{G}^j and simply write $\mathcal{F}(\tilde{\theta}^i | \omega^k)$ where clear. Given \mathcal{F} , we can calculate, for each $\theta^i \in \Theta$, $\omega^k \in \Omega$, the difference $d(\theta^i, \omega^k)$ between the empirical frequency and theoretical probability of observing θ^i given ω^k , as

$$d(\theta^i, \omega^k) = |\mathcal{F}(\theta^i | \omega^k) - f(\theta^i | \omega^k)|$$

Conditional on being in review phase \mathcal{G}^j , the Mediator performs the following tests in period $\mathcal{G}(j) + T - 1$. At this time, the Mediator calculates the realized values of $d(\theta^i, \omega^k)$ for all $\theta^i \in \Theta$, and all $\omega^k \in \Omega$. The sender *passes the review phase* \mathcal{G}^j if

$$d(\theta^i, \omega^k) < \chi,$$

for all $\theta^i \in \Theta$, and all $\omega^k \in \Omega$ satisfying $|\kappa(\omega^k, \mathcal{G}^j)| > \underline{\kappa}$, for some fixed parameter $\underline{\kappa} \in \mathbb{N}_{>0}$, where $\chi > 0$ is a fixed parameter of the mechanism. Otherwise, he *fails the review*. If the sender passes, then the Mediator sets $\mathcal{B}(j) = 0$. Otherwise, $\mathcal{B}(j) = \beta T$.

On the Connection between Mediated Communication and Review Aggregation

Notice the differences between mediated communication and review aggregation. The review aggregator only tracks the sender's interactions with receivers and based on this, determines the history observable to receivers. In addition to these tests, the mediator interacts directly with the sender by asking for reports, $\tilde{\theta}_t$, and based on this implements a target reporting strategy r to each receiver. Thus, the mediator is more intrusive than the Aggregator. In particular, this implies the set of deviations available to the sender is effectively lower under mediated communication. Moreover review aggregation mechanisms are routinely used in practice in many economic settings, whereas mediated communication of the above form is less prevalent.

Nonetheless, we present the result in this section for mediated communication for the following reasons: first, it allows us to illustrate that the results hold in a more traditional mechanism design setting. Second, by basing tests on truth-telling, we are able to pin down the sender's optimal strategy more finely and establish that the frequency of 'lies' vanishes in the limit. This formalizes our claim in the main text that we obtain a large increase in attainable payoffs by allowing only a very small violation of incentive compatibility. Third, given the reduced set of available deviations, the sender's compliance with statistical tests can be ensured at a lower δ – implying that mediated communication can offer efficiency gains. Finally, it clarifies the exposition of the proof.

However, we note that the general result presented below can be replicated for the case of a review aggregator under Assumptions 3 and 4. In particular, the argument in the proof of Theorem 3 extends naturally, albeit at some notational burden.

The Main Result

With the correct choice of Review Mechanism and a patient enough sender, it turns out we can recover payoffs to the sender arbitrarily close to his optimal commitment payoff, even when the sender's private information is not verifiable ex post:

Theorem 4. *Suppose Assumptions 3 and 4 hold. For any $\varepsilon > 0$, there exist parameters $\Gamma_{\mathcal{R}}$ and $\underline{\delta} < 1$ such that for all $1 > \delta \geq \underline{\delta}$ there is an equilibrium of review mechanism $\Gamma_{\mathcal{R}}$ in which the sender's normalized discounted average utility is at least $\hat{v}(\mu_0) - \varepsilon$:*

$$(1 - \delta) \mathbb{E} \left[\sum_{\tau=t}^{\infty} \delta^{\tau} u_S(a_{\tau}) \mid \tilde{\theta}^*, \rho^*; \Gamma_{\mathcal{R}} \right] \geq \hat{v}(\mu_0) - \varepsilon.$$

where $\mathbb{E} \left[\cdot \mid \tilde{\theta}^*, \rho^*; \Gamma_{\mathcal{R}} \right]$ denotes expectations taken with respect to equilibrium play of $\Gamma_{\mathcal{R}}$.

Proof of Theorem 4

Proof. In arbitrary review phase \mathcal{G}_j , denote realized sequences of the sender reports, types and receiver signals respectively by

$$\begin{aligned} (\tilde{\theta})^{\mathcal{G}(j)+T-1|\mathcal{G}(j)} &= \left(\tilde{\theta}_{\mathcal{G}(j)}, \tilde{\theta}_{\mathcal{G}(j)+1}, \dots, \tilde{\theta}_{\mathcal{G}(j)+T-1} \right), \\ (\theta)^{\mathcal{G}(j)+T-1|\mathcal{G}(j)} &= \left(\theta_{\mathcal{G}(j)}, \theta_{\mathcal{G}(j)+1}, \dots, \theta_{\mathcal{G}(j)+T-1} \right), \end{aligned}$$

and

$$(\omega)^{\mathcal{G}(j)+T-1|\mathcal{G}(j)} = \left(\omega_{\mathcal{G}(j)}, \omega_{\mathcal{G}(j)+1}, \dots, \omega_{\mathcal{G}(j)+T-1} \right).$$

Similarly, in punishment phases \mathcal{B}_j define

$$\begin{aligned} (\tilde{\theta})^{\mathcal{G}(j)+(1+\beta)T|\mathcal{G}(j)+T}, \\ (\theta)^{\mathcal{G}(j)+(1+\beta)T|\mathcal{G}(j)+T}, \\ (\omega)^{\mathcal{G}(j)+(1+\beta)T|\mathcal{G}(j)+T}. \end{aligned}$$

Due to the recursive nature of the equilibrium we establish below, the same arguments will hold across all \mathcal{G}_j , \mathcal{B}_j , $j = 1, 2, \dots$, respectively, and therefore it is sufficient to establish arguments for $j = 1$. To ease notation, we refer to the above sequences as $(\tilde{\theta})_{\mathcal{G}}^{T-1} := (\tilde{\theta})^{\mathcal{G}(1)+T-1|\mathcal{G}(1)}$, $(\theta)_{\mathcal{G}}^{T-1} := (\theta)^{\mathcal{G}(1)+T-1|\mathcal{G}(1)}$ and $(\omega)_{\mathcal{G}}^{T-1} := (\omega)^{\mathcal{G}(1)+T-1|\mathcal{G}(1)}$ in phase \mathcal{G}_0 and $(\tilde{\theta})_{\mathcal{B}}^{(1+\beta)T}$, $(\theta)_{\mathcal{B}}^{(1+\beta)T}$, $(\omega)_{\mathcal{B}}^{(1+\beta)T}$ in \mathcal{B}_0 . For any $(\theta^i, \omega^j) \in \Theta \times \Omega$, we define the empirical frequency of the joint observation (θ_t^i, ω_t^j) given sequences $(\theta)_{\mathcal{G}}^{T-1}$ and $(\omega)_{\mathcal{G}}^{T-1}$ as $\mathcal{F}_T(\theta^i, \omega^j) := \frac{\sum_{t=0}^{T-1} \mathbf{1}(\theta_t^i, \omega_t^j)}{T}$. Similarly, given $(\tilde{\theta})_{\mathcal{G}}^{T-1}$, $(\omega)_{\mathcal{G}}^{T-1}$, let the empirical frequency of joint observation $(\tilde{\theta}_t = \theta^i, \omega_t^j)$ be $\tilde{\mathcal{F}}_T(\tilde{\theta} = \theta^i, \omega^j) := \frac{\sum_{t=0}^{T-1} \mathbf{1}(\tilde{\theta}_t = \theta^i, \omega_t^j)}{T}$. Finally, given

sequences $(\tilde{\theta})_{\mathcal{G}}^{T-1}$, $(\theta)_{\mathcal{G}}^{T-1}$, let $b(\tilde{\theta} = \theta^i | \theta^k)$, $i, k \in \{1, 2, \dots, N\}$, be the frequency of reports $\tilde{\theta} = \theta^i$ when S observes θ^k ,

$$b(\tilde{\theta} = \theta^i | \theta^k) := \begin{cases} \frac{\sum_{t=0}^{T-1} \mathbf{1}(\tilde{\theta}_t = \theta^i, \theta_t^k)}{\sum_{t=0}^{T-1} \mathbf{1}(\theta_t^k)} & , \text{ if } \sum_{t=1}^T \mathbf{1}(\theta^k) > 0 \\ 0 & , \text{ otherwise.} \end{cases}$$

Consider review phase \mathcal{G}_0 and fix $\epsilon_T, \xi_T > 0$. We first show that there exist $\hat{\chi}, \hat{T}$ such that (i) if the sender adopts a truth-telling strategy within review phase \mathcal{G}_0 , he passes the review with probability at least $\Pr\left(\bigcap_{i,j \in \{1,2,\dots,N\}} \left\{ \left| \hat{\mathcal{F}}_T(\tilde{\theta} = \theta^i, \omega^j) - f(\theta^i, \omega^j) \right| \leq \chi \right\}\right) \geq 1 - \epsilon_{\bar{T}}$;

(ii) $\Pr\left(\bigcap_{i,j \in \{1,2,\dots,N\}} \left\{ |\mathcal{F}_T(\theta^i) - \mu_0^i| \leq \chi \right\}\right) \geq 1 - \epsilon_{\bar{T}}$; (iii) for any sequences $(\tilde{\theta})_{\mathcal{G}}^{T-1}$, $(\theta)_{\mathcal{G}}^{T-1}$ such that $|\mathcal{F}_T(\theta^i) - \mu_0^i| \leq \chi$, $i = 1, 2, \dots, N$, if $\left| \left(b(\tilde{\theta} = \theta^k | \theta^i) \right)_{k=1}^N - \mathbf{e}_{\theta^i} \right| > \xi_T$, then S fails the review with probability at least $1 - \epsilon_{\bar{T}}$, where $\mathbf{e}_{\theta^i} = (0, 0, \dots, 1, \dots, 0)$ is an $N \times 1$ vector whose i^{th} row is 1, and all others are 0.

To establish (i), choose $\chi(T) = \frac{\psi}{T^y}$, $\psi > 0$, $0 < y < \frac{1}{2}$, for suppose that S uses the truthful strategy $\hat{\theta}_t((\tilde{\theta})^{t-1}, (\theta)^{t-1}, \theta_t) = \theta_t$. Under the truthful strategy, $(\tilde{\theta}, \omega)$ is a sequence of independent Bernoulli trials. By Chebyshev's inequality,

$$\Pr\left(\left| \hat{\mathcal{F}}_T(\hat{\theta} = \theta^i, \omega^j) - f(\theta^i, \omega^j) \right| \leq \chi(T)\right) \geq 1 - \frac{\psi^{-2}}{4T^{1-2y}}$$

Letting $A_{ij} := \left\{ (\tilde{\theta}, \theta, \omega)_{\mathcal{G}}^{T-1} : \left| \hat{\mathcal{F}}_T(\hat{\theta} = \theta^i, \omega^j) - f(\theta^i, \omega^j) \right| \leq \chi(T) \right\}$, we can write

$$\Pr\left(\bigcap_{i,j \in \{1,2,\dots,N\}} A_{ij}\right) \geq 1 - \sum_{i,j \in \{1,2,\dots,N\}} \Pr(\bar{A}_{ij}) = 1 - \frac{N^2 \psi^{-2}}{4T^{1-2y}}$$

Choosing T' , $\chi(T')$ for any $T' \geq \bar{T}$, where $\frac{N^2 \psi^{-2}}{4T'^{1-2y}} \leq \epsilon_{\bar{T}}$, establishes part (i). Similarly, claim (ii) can be shown to follow from Chebyshev's inequality for any T' , $\chi(T')$ such that $T' \geq \bar{T}$.

To establish (iii), fix $T' \geq \bar{T}$ and sequences $(\tilde{\theta})_{\mathcal{G}}^{T'-1}$, $(\theta)_{\mathcal{G}}^{T'-1}$ such that $|\mathcal{F}_T(\theta^k) - \mu_0^k| \leq \chi(T')$, $\forall k \in \{1, 2, \dots, N\}$. Consider the empirical frequency $\tilde{\mathcal{F}}_{T'}(\tilde{\theta} = \theta^i, \omega^j)$, conditional on $(\tilde{\theta})_{\mathcal{G}}^{T'-1}$, $(\theta)_{\mathcal{G}}^{T'-1}$.

At each $t \in \{1, 2, \dots, T'\}$, the event $(\tilde{\theta}_t = \theta^i, \omega_t^j)$ is a Bernoulli trial with probability of success, conditional on $(\theta)_{\mathcal{G}}^{T'-1}$

$$\begin{cases} f(\omega^j | \theta^k) & , \text{ if } \tilde{\theta} = \theta^i, \theta_t = \theta^k, k \in \{1, 2, \dots, N\} \\ 0 & , \text{ otherwise.} \end{cases}$$

Moreover, since (given θ_t) ω_t is conditionally independent of all $\theta_\tau, \omega_\tau, \tau \neq t$, we have a sequence of independent Bernoulli trials.

Calculating the expectation of $\tilde{\mathcal{F}}_{T'}(\tilde{\theta} = \theta^i, \omega^j)$, given $(\tilde{\theta})_{\mathcal{G}}^{T'-1}, (\theta)_{\mathcal{G}}^{T'-1}$, we have

$$\mathbb{E} \left[\tilde{\mathcal{F}}_{T'}(\tilde{\theta} = \theta^i, \omega^j) \mid (\tilde{\theta}, \theta)_{\mathcal{G}}^{T'-1} \right] = \sum_{\theta^k \in \Theta} b(\tilde{\theta} = \theta^i \mid \theta^k) \mathcal{F}_{T'}(\theta^k) f(\omega^j \mid \theta^k)$$

Since $\left\{ (\tilde{\theta}_t = \theta^i, \omega_t^j) \right\}_{t=0}^{T'}$ is a sequence of independent Bernoulli trials, it follows from Chebyshev's inequality that

$$\Pr \left(\left| \tilde{\mathcal{F}}_{T'}(\tilde{\theta} = \theta^i, \omega^j) - \mathbb{E} \left[\tilde{\mathcal{F}}_{T'}(\tilde{\theta} = \theta^i, \omega^j) \mid (\tilde{\theta}, \theta)_{\mathcal{G}}^{T'-1} \right] \right| \leq \chi(T') \mid (\tilde{\theta}, \theta)_{\mathcal{G}}^{T'-1} \right) \geq 1 - \frac{\psi^{-2}}{4(T')^{1-2y}} \quad (26)$$

Now, since $|\mathcal{F}_{T'}(\theta^k) - \mu_0^k| \leq \chi(T')$, and by definition

$$0 \leq \sum_{\theta^k \in \Theta} b(\tilde{\theta} = \theta^i \mid \theta^k) f(\omega^j \mid \theta^k) < \sum_{\theta^k \in \Theta} b(\tilde{\theta} = \theta^i \mid \theta^k) \leq N,$$

we can bound

$$\left| \mathbb{E} \left[\tilde{\mathcal{F}}_{T'}(\tilde{\theta} = \theta^i, \omega^j) \mid (\tilde{\theta}, \theta)_{\mathcal{G}}^{T'-1} \right] - \sum_{\theta^k \in \Theta} b(\tilde{\theta} = \theta^i \mid \theta^k) f(\theta^k, \omega^j) \right| \leq N\chi(T'),$$

which follows after recalling that $f(\theta^k, \omega^j) = \mu_0^k f(\omega^j \mid \theta^k), \forall \theta^k$. Using the triangle inequality and (26), we can bound $\Pr \left(\tilde{A}_{i,j} \mid (\tilde{\theta}, \theta)_{\mathcal{G}}^{T'-1} \right) \geq 1 - \frac{\psi^{-2}}{4(T')^{1-2y}}$, where

$$\tilde{A}_{i,j} := \left\{ (\tilde{\theta}, \theta, \omega)_{\mathcal{G}}^{T'-1} : \left| \tilde{\mathcal{F}}_{T'}(\theta^i, \omega^j) - \sum_{\theta^k \in \Theta} b(\theta^i \mid \theta^k) f(\theta^k, \omega^j) \right| \leq (N+1)\chi(T') \right\},$$

It is easy to see that this implies $\Pr \left(\cup_{i,j \in \{1,2,\dots,N\}} \tilde{A}_{i,j} \mid (\tilde{\theta}, \theta)_{\mathcal{G}}^{T'-1} \right) \geq 1 - \epsilon_{T'}$.

The events $\tilde{A}_{i,j}, i, j \in \{1, 2, \dots, N\}$, can be written using the system of inequalities

$$\mathbf{F}_{\theta, \omega} \cdot \mathbf{b}_i - (N+1)\chi(T') \cdot \mathbf{1} \leq \tilde{\mathcal{F}}_{T'} \leq \mathbf{F}_{\theta, \omega} \cdot \mathbf{b}_i + (N+1)\chi(T') \cdot \mathbf{1} \quad (27)$$

where

$$\mathbf{F}_{\theta, \omega} = \begin{pmatrix} f(\theta^1, \omega^1) & \cdots & f(\theta^N, \omega^1) \\ \vdots & \ddots & \vdots \\ f(\theta^N, \omega^1) & \cdots & f(\theta^N, \omega^N) \end{pmatrix}, \tilde{\mathcal{F}}_{T'} = \begin{pmatrix} \tilde{\mathcal{F}}_{T'}(\theta^i, \omega^1) \\ \vdots \\ \tilde{\mathcal{F}}_{T'}(\theta^i, \omega^N) \end{pmatrix},$$

and

$$\mathbf{b}_i = \begin{pmatrix} b(\theta^i | \theta^1) \\ \vdots \\ b(\theta^i | \theta^N) \end{pmatrix}, \mathbf{1} = \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix}.$$

Assumption 3 implies that $\mathbf{F}_{\theta,\omega}$ is invertible. Thus, the linear function $\mathbf{F}_{\theta,\omega} : \mathbb{R}^N \rightarrow \mathbb{R}^N$ is continuous and injective. It follows that for any $\xi > 0$, we can find v_ξ such that if $|\mathbf{b}_i - \mathbf{e}_{\theta_i}| > \xi$, then $|\mathbf{F}_{\theta,\omega} \cdot \mathbf{b}_i - \mathbf{F}_{\theta,\omega} \cdot \mathbf{e}_{\theta_i}| = |\mathbf{F}_{\theta,\omega} \cdot \mathbf{b}_i - \mathbf{F}_{\theta,\omega} \cdot \mathbf{e}_{\theta_i}| > v_\xi$, where $v_\xi \rightarrow 0$ as $\xi \rightarrow 0$. Given this v_ξ , we can select T_ξ such that $2(N+1)\chi(T_\xi) \leq v_\xi$. Thus, for at least one $i \in \{1, 2, \dots, N\}$, we have

$$v_\xi - (N+1)\chi(T_\xi) > 0$$

$$\begin{aligned} \left| \tilde{\mathcal{F}}_{T_\xi}(\theta^i, \omega^j) - f(\theta^i, \omega^j) \right| &\geq \left| \mathbf{F}_{\theta,\omega}^i \cdot \mathbf{b}_i - \mathbf{F}_{\theta,\omega}^i \cdot \mathbf{e}_{\theta_i} \right| - \left| \tilde{\mathcal{F}}_{T'}(\theta^i, \omega^j) - \mathbf{F}_{\theta,\omega}^i \cdot \mathbf{b}_i \right| \\ &\geq v_\xi - (N+1)\chi(T_\xi) \\ &\geq (N+1)\chi(T_\xi) \end{aligned}$$

Applying Chebyshev's inequality again, we have found a T_ξ such that for any $|\mathbf{b}_i - \mathbf{e}_{\theta_i}| > \xi$, there is $i \in N$ such that

$$\begin{aligned} \Pr \left(\left| \tilde{\mathcal{F}}_{T'}(\theta^i, \omega^j) - f(\theta^i, \omega^j) \right| > \chi(T_\xi) \right) &\geq \\ \Pr \left(\left| \tilde{\mathcal{F}}_{T'}(\theta^i, \omega^j) - \mathbf{F}_{\theta,\omega}^i \cdot \mathbf{b}_i \right| \leq (N+1)\chi(T_\xi) \right) &\geq \\ &1 - \frac{\epsilon_{T_\xi}}{N^2} \end{aligned}$$

Selecting $\hat{T} = \max\{\bar{T}, T_\xi\}$, $\hat{\chi} = \chi(\hat{T})$ establishes claim (iii).

Next, fix $\hat{\chi}$, \hat{T} , and some Mediator report function r such that $r(\mathcal{G}, \theta) = \sigma : \Theta \rightarrow \Delta\Theta$, with support $\tilde{\mu} \in \{\mu_1, \dots, \mu_N\}$, and $r(\mathcal{B}, \theta) = \underline{\sigma} : \Theta \rightarrow \Delta\Theta$ with support $\tilde{\mu} \in \{\underline{\mu}_1, \dots, \underline{\mu}_N\}$, where $\underline{\sigma}$ is the sender's reporting strategy in the worst equilibrium of the stage game in Section 2. Suppose further that receiver posterior beliefs satisfy, for all $t = 0, 1, \dots, T-1$,

$$\hat{\mu}(\tilde{\mu}_t, \mathcal{P}) = \begin{cases} \hat{\mu}_i & , \text{ if } \mathcal{P} = \mathcal{G}, \tilde{\mu} = \mu_i \\ \underline{\mu}_i & , \text{ if } \mathcal{P} = \mathcal{B}, \tilde{\mu} = \mu_i \end{cases}$$

where $\hat{\mu} \in \Delta\Theta$.⁵¹ Let Q_t denote the set of all pairs of sequences $q^t = \left((\tilde{\theta})^{t-1}, (\theta)^t \right)$, and $Q = \cup_{t \in \{0, 1, \dots, T-1\}} Q_t$. With some notational abuse, let $\tilde{\Theta}^Q$ denote the set of feasible reporting *plans* within phase \mathcal{G} , $\tilde{\theta}^Q := \left(\tilde{\theta}_t : \Theta^{t-1} \times \Theta^t \rightarrow \Theta \right)_{t=0}^{T-1}$. Given a reporting plan

⁵¹We show below that in equilibrium, receiver beliefs are indeed invariant to the calendar time t of the receiver at date, conditional on report phase \mathcal{G} .

$\tilde{\theta}_t(q^t)$ and report function σ , we write the induced lottery over $\{\mu_1, \dots, \mu_N\}$, conditional on $(\tilde{\theta})^{t-1}, (\theta)^{t-1}$ as $\lambda_i(q^t) := \Pr(\tilde{\mu} = \mu_i | \tilde{\theta}^{t-1}, \theta^{t-1})$. Associated with this lottery, define the induced lottery over receiver posteriors as $\hat{\lambda}$, where $\hat{\lambda}_i := \Pr(\hat{\mu} = \hat{\mu}_i | \tilde{\theta}^{t-1}, \theta^{t-1})$. Note that in any PBE, $\hat{\lambda} \in \Lambda(\mu_0)$.

Clearly, in any \mathcal{B} phase, truthful reporting by the sender can be sustained as part of a PBE, since Mediator's reports (conditional on $\tilde{\theta}$) and receiver beliefs are specified as those in the PBE of the worst stage game. For convenience, we normalize $\sum_{i=1}^N \lambda_j v(\mu_j) = 0$.

Consider the sender's best response in review phase, \mathcal{G} . Letting $\tilde{\varphi}$ denote the probability that S fails the review at the end of period $T - 1$ when following reporting plan $\tilde{\theta}^Q$, we can write the sender's payoffs at the beginning of phase \mathcal{G} recursively as

$$\mathcal{V}_{\mathcal{G}} = \max_{\tilde{\theta}^Q \in \Theta^Q} \sum_{t=0}^{T-1} \delta^t \mathbb{E} \left[\sum_{i=1}^N \hat{\lambda}_i(\tilde{\theta}^t(q^t)) v(\hat{\mu}_{t,i}) \right] + \delta^T [(1 - \tilde{\varphi}) + \tilde{\varphi} \delta^{\beta T}] \mathcal{V}_{\mathcal{G}} \quad (28)$$

(28) is a standard dynamic programming problem with a finite set of states - therefore, a stationary optimal strategy for the sender exists. Letting the best response function be $\tilde{\theta}^t = \bar{\theta}^t$, we can equivalently write (28) as⁵²

$$(1 - \delta) \mathcal{V}_{\mathcal{G}} = (1 - \delta) \frac{\sum_{t=0}^{T-1} \delta^t \mathbb{E} \left[\sum_{i=1}^N \hat{\lambda}_i(\bar{\theta}^t(q^t)) v(\hat{\mu}_{t,i}) \right]}{1 - (1 - \tilde{\varphi}) \delta^T - \tilde{\varphi} \delta^{(1+\beta)T}} \quad (29)$$

or $(1 - \delta) \mathcal{V}_{\mathcal{G}} =$

$$\frac{(1 - \delta)}{1 - \delta^T} \left[\sum_{t=0}^{T-1} \delta^t \mathbb{E} \left[\sum_{i=1}^N \hat{\lambda}_i(\bar{\theta}^t(q^t)) v(\hat{\mu}_{t,i}) \right] - \delta^T (1 - \delta^{\beta T}) \tilde{\varphi} (1 - \delta) \mathcal{V}_{\mathcal{G}} \right] \quad (30)$$

We now argue that, for any $\underline{\varphi} \in (0, 1)$ and $\varepsilon > 0$ and $\eta > 0$, there exist $\bar{\chi}, \bar{T}, r, \bar{\beta}, \delta_{\bar{T}}$, such that if receiver posteriors given Mediator message $\tilde{\mu}_i$ satisfy $|\hat{\mu}_i - \tilde{\mu}_i| < \eta$ for all $i \in \{1, 2, \dots, N\}$, then (i) the sender's optimal strategy involves $\tilde{\varphi} \leq \underline{\varphi}$, and (ii) $(1 - \delta) \mathcal{V}_{\mathcal{G}} \geq \hat{v}(\mu_0) - \varepsilon$, for all $\delta \geq \delta_{\bar{T}}$.

Fix a Mediator report function $r(\mathcal{G}, \tilde{\theta}_t) = \sigma' : \Theta \rightarrow \Delta\Theta$, with support $\{\mu'_1, \dots, \mu'_N\}$ and induced lottery λ' , where λ' satisfies $\lambda' \in X(\lambda^{**})$ for some optimal experiment under commitment, $\lambda^{**} \in \Lambda(\mu_0)$, and

$$\left| \sum_{i=1}^N \lambda'_i(\theta_t) v(\mu'_i) - \hat{v}(\mu_0) \right| \leq \frac{\varepsilon}{4}$$

. Such a report function exists under Assumption 4. Suppose further that receiver posterior beliefs, $\hat{\mu}'(\tilde{\mu})$, satisfy $|\hat{\mu}'_i - \tilde{\mu}_i| < \eta'$, for all $i \in \{1, 2, \dots, N\}$, where η' is small

⁵²We suppress explicit dependence of $\bar{\theta}$ on $\hat{\chi}, \bar{T}, r, \hat{\mu}, \delta, \beta$ for notational ease.

enough that

$$\left| \sum_{i=1}^N \lambda'_i(\theta_t) v(\mu'_i) - \sum_j \lambda'_i(\theta_t) v(\hat{\mu}_i) \right| \leq \frac{\varepsilon}{4}$$

for all $\hat{\mu}'_i \in N_{\eta'}(\tilde{\mu}_i)$, $i = 1, \dots, N$.

Suppose that $(1 - \delta) \mathcal{V}_{\mathcal{G}} \geq \hat{v}(\mu_0) - \frac{3\varepsilon}{4}$ for all $\delta \geq \delta_1$ (verified below). Taking limits of equation (30) as $\delta \rightarrow 1$ for any reporting plan $\tilde{\theta}_t(q^t)$ in phase \mathcal{G} that involves a probability of failing the review of at least $\tilde{\varphi}$, we can write an upper bound the sender's payoff as

$$\begin{aligned} \lim_{\delta \rightarrow 1} (1 - \delta) \mathcal{V}_{\mathcal{G}} &\leq \bar{v} - \beta \tilde{\varphi} \lim_{\delta \rightarrow 1} (1 - \delta) \mathcal{V}_{\mathcal{G}} \\ &\leq \bar{v} - \beta \tilde{\varphi} \left(\hat{v}(\mu_0) - \frac{3\varepsilon}{4} \right) \end{aligned}$$

where $\bar{v} := \max \{v(\mu_i)\}_{i=1}^N$. Clearly, given a $\bar{\varphi} > 0$, we can find a $\bar{\beta} \in \mathbb{N}_{>0}$ such that $\bar{v} - \bar{\beta} \bar{\varphi} \left(\hat{v}(\mu_0) - \frac{3\varepsilon}{4} \right) \leq \hat{v}(\mu_0) - \frac{3\varepsilon}{4}$ - a contradiction to the lower bound, $(1 - \delta) \mathcal{V}_{\mathcal{G}} \geq \hat{v}(\mu_0) - \frac{3\varepsilon}{4}$. Thus, for any $\beta \geq \bar{\beta}$, it cannot be optimal (in the limit) for the sender to use a reporting plan $\tilde{\theta}_t(q^t)$ in phase \mathcal{G} that $\tilde{\varphi} > \bar{\varphi}$. We now establish that such a lower bound can be imposed, given $\bar{\beta}$. Fixing $\bar{\beta} > 0$, note that we can bound the sender's limit payoff, $\lim_{\delta \rightarrow 1} (1 - \delta) \mathcal{V}_{\mathcal{G}}$, as $\delta \rightarrow 1$ in each \mathcal{G} -phase by the sender's limit payoff from truthful reporting. From equation (29),

$$\begin{aligned} \lim_{\delta \rightarrow 1} (1 - \delta) \mathcal{V}_{\mathcal{G}} &\geq \frac{T \cdot \mathbb{E} \left[\sum_{i=1}^N \hat{\lambda}_i(\theta_t) v(\hat{\mu}_{t,i}) \right]}{T + \epsilon_T \bar{\beta} T} \\ &\geq \frac{\hat{v}(\mu_0) - \frac{\varepsilon}{2}}{1 + \epsilon_T \bar{\beta}} \end{aligned}$$

By choosing \bar{T} sufficiently high, and $\bar{\chi} = \chi(\bar{T})$ we can set $\epsilon_{\bar{T}}$ small enough that $\frac{\hat{v}(\mu_0) - \frac{\varepsilon}{2}}{1 + \epsilon_T \bar{\beta}} \geq \hat{v}(\mu_0) - \frac{3\varepsilon}{4}$. Denoting by $\tilde{\mathcal{V}}_{\mathcal{G}}$ the discounted payoff from an arbitrary recursive strategy of using reporting plan $\tilde{\theta}_t(q^t)$ in phase \mathcal{G} , and $\tilde{\theta}_t = \theta_t$ in phase \mathcal{B} . Note that there are only finitely many such choices of recursive strategy, because each phase has finitely many nodes and messages. Given \bar{T} , we can therefore select $\delta_{\bar{T}} < 1$ such that $\left| (1 - \delta) \tilde{\mathcal{V}}_{\mathcal{G}} - \lim (1 - \delta) \tilde{\mathcal{V}}_{\mathcal{G}} \right| < \frac{\varepsilon}{4}$ uniformly across all recursive strategies. In particular, since the optimal strategy is recursive, this implies that $(1 - \delta) \mathcal{V}_{\mathcal{G}} \geq \hat{v}(\mu_0) - \varepsilon$, for all $\delta \geq \delta_{\bar{T}}$ - as required.

Given that A is countable, Assumption 4 ensures there exists an $\eta'' > 0$ such that $a^*(\hat{\mu}'_i) = a^*(\tilde{\mu}_i)$, for all $|\hat{\mu}'_i - \tilde{\mu}_i| < \eta''$. Thus, to complete the proof of existence of an equilibrium with desired properties, we need only show there exists $\bar{\varphi}$ such that receiver beliefs are guaranteed to satisfy $|\hat{\mu}'_i - \tilde{\mu}_i| < \eta''$ in equilibrium - the sender's and receivers' optimal strategies are then mutual best responses, and beliefs are correct.

For any $\bar{\varphi} \in (0, 1)$, we have shown there exist $\bar{\chi}, \bar{T}, r, \bar{\beta}, \delta_{\bar{T}}$ such that $\tilde{\varphi} \leq \bar{\varphi}$, and (ii)

$(1 - \delta) \mathcal{V}_{\mathcal{G}} \geq \hat{v}(\mu_0) - \varepsilon$, for all $\delta \geq \delta_{\bar{T}}$. Additionally, for any $\epsilon_{T^*} \leq \epsilon_{\bar{T}}$, $\xi_{T^*} > 0$ we can choose $\chi(T^*)$, T^* , where $T^* = \max\{\bar{T}, \hat{T}\}$, to ensure that sequences $(\theta)_{\mathcal{G}}^{T-1}$ satisfying $|\mathcal{F}_T(\theta^i) - \mu_0^i| \leq \chi(T^*)$, for all $i = 1, 2, \dots, N$, occur with probability greater than $1 - \epsilon_{T^*}$ in review phase \mathcal{G} (claim (i)); conditional on any such sequence, a reporting strategy that involves $|\mathbf{b}_i - \mathbf{e}_{\theta^i}| > \xi_{T^*}$ fails the review with probability at least $1 - \epsilon_{\bar{T}}$ (claim (iii)).

Combining these observations, we can bound the probability that the sender's optimal reporting plan $(\bar{\theta}_t(q_t))_{t=0}^{T-1}$ involves $|\mathbf{b}_i - \mathbf{e}_{\theta^i}| > \xi$ at some history $(\bar{\theta}(q_t))^{T-1}$, $(\theta)^{T-1}$, for some $i = 1, 2, \dots, N$, as follows. By definition

$$\begin{aligned} \Pr(|\mathbf{b}_i - \mathbf{e}_{\theta^i}| > \xi_{T^*}) &= \Pr(\cap_i \{|\mathcal{F}_{T^*}(\theta) - \mu_0^i| \leq \chi^*\}, |\mathbf{b}_i - \mathbf{e}_{\theta^i}| > \xi_{T^*}) \\ &+ \Pr(\cup_i \{|\mathcal{F}_{T^*}(\theta) - \mu_0^i| > \chi^*\}, |\mathbf{b}_i - \mathbf{e}_{\theta^i}| > \xi_{T^*}) \end{aligned}$$

for any $i = 1, 2, \dots, N$. For all such i , we can bound each term on the right hand-side of the inequality respectively by

$$\begin{aligned} \Pr(\cap_i \{|\mathcal{F}_{T^*}(\theta) - \mu_0^i| \leq \chi^*\}, |\mathbf{b}_i - \mathbf{e}_{\theta^i}| > \xi_{T^*}) (1 - \epsilon_{T^*}) &\leq \\ \Pr(S \text{ fails review at } T^* - 1) &\leq \bar{\varphi} \end{aligned}$$

and

$$\Pr(\cup_i \{|\mathcal{F}_{T^*}(\theta) - \mu_0^i| > \chi^*\}, |\mathbf{b}_i - \mathbf{e}_{\theta^i}| > \xi_{T^*}) \leq \Pr(\cup_i |\mathcal{F}_{T^*}(\theta) - \mu_0^i| > \chi^*) \leq \epsilon_{T^*}$$

Thus, $\Pr(|\mathbf{b}_i - \mathbf{e}_{\theta^i}| > \xi_{T^*}) \leq \frac{\bar{\varphi}}{1 - \epsilon_{T^*}} + \epsilon_{T^*} \rightarrow 0$, $\xi_{T^*} \rightarrow 0$ as $T^* \rightarrow \infty$. Let $\mathcal{T} \in 2^{\{0, 1, \dots, T^*-1\}}$ denote an arbitrary subset of $\{0, 1, \dots, T^* - 1\}$ and for any possible history $(\tilde{\theta})_{\mathcal{G}}^{T-1}$, $(\theta)_{\mathcal{G}}^{T-1}$ define $b^{\mathcal{T}}(\theta_i, \theta_j) = \frac{\sum_{t \in \mathcal{T}} \mathbf{1}(\tilde{\theta}_t = \theta_i) \cdot \mathbf{1}(\theta_t = \theta_j)}{\sum_{t \in \mathcal{T}} \mathbf{1}(\theta_t = \theta_j)}$. Similarly, let $\mathbf{b}_i^{\mathcal{T}} := (b^{\mathcal{T}}(\theta_i, \theta_j))_{j=1}^N$. On any history $(\tilde{\theta})_{\mathcal{G}}^{T-1}$, $(\theta)_{\mathcal{G}}^{T-1}$ define \mathcal{T}_L^{ξ} as the solution to

$$T_L(\xi) := \max_{\mathcal{T} \in 2^{\{0, 1, \dots, T^*-1\}}} |\mathcal{T}| \quad (31)$$

s.t.

$$|\mathbf{b}_i^{\mathcal{T}} - \mathbf{e}_{\theta^i}| \geq \xi^{\frac{1}{2}}$$

where $|\mathcal{T}|$ denotes the cardinality of set \mathcal{T} . Since problem (31) is a finite choice problem, it has a well-defined solution. It is easy to show arithmetically that for any history $(\tilde{\theta})_{\mathcal{G}}^{T-1}$, $(\theta)_{\mathcal{G}}^{T-1}$ such that $|\mathbf{b}_i - \mathbf{e}_{\theta^i}| \leq \xi$, we have $\frac{T_L(\xi)}{T} \leq \xi^{\frac{1}{2}}$.

Now, consider receiver R_t 's inference problem, given observations $(\mathcal{G}, \hat{\mu})$. Given uniform permutation of receivers in any phase of the review mechanism, R_t 's beliefs do not

depend on her index t . Letting

$$(|\mathbf{B} - \mathbf{E}| \leq \xi) := \left\{ \left(\tilde{\theta}(q^t), \theta \right)_{\mathcal{G}}^{T-1} : |\mathbf{b}_i - \mathbf{e}_{\theta_i}| \leq \xi, \forall i \in \{1, \dots, N\} \right\},$$

we can write posterior beliefs given message $\tilde{\mu}$ as

$$\begin{aligned} \Pr(\theta_{\tilde{\pi}(t)} = \theta_i \mid \mathcal{G}, \tilde{\mu}) &= \Pr(\theta_{\tilde{\pi}(t)} = \theta_i, |\mathbf{B} - \mathbf{E}| \leq \xi^{\frac{1}{2}}, \tilde{\pi}(t) \notin \mathcal{T}_L^\xi \mid \mathcal{G}, \tilde{\mu}) \\ &+ \Pr(\theta_{\tilde{\pi}(t)} = \theta_i, |\mathbf{B} - \mathbf{E}| \leq \xi^{\frac{1}{2}}, \tilde{\pi}(t) \in \mathcal{T}_L^\xi \mid \mathcal{G}, \tilde{\mu}) \\ &+ \Pr(\theta_{\tilde{\pi}(t)} = \theta_i, |\mathbf{B} - \mathbf{E}| > \xi^{\frac{1}{2}} \mid \mathcal{G}, \tilde{\mu}) \end{aligned} \quad (32)$$

The final two terms can be bounded respectively by

$$\begin{aligned} &\Pr(\theta_{\tilde{\pi}(t)} = \theta_i, |\mathbf{B} - \mathbf{E}| \leq \xi^{\frac{1}{2}}, \tilde{\pi}(t) \in \mathcal{T}_L^\xi \mid \mathcal{G}, \tilde{\mu}) \\ &\leq \Pr(|\mathbf{B} - \mathbf{E}| \leq \xi^{\frac{1}{2}}, \tilde{\pi}(t) \in \mathcal{T}_L^\xi \mid \mathcal{G}, \tilde{\mu}) \\ &\leq \Pr(\tilde{\pi}(t) \in \mathcal{T}_L^\xi \mid |\mathbf{B} - \mathbf{E}| \leq \xi^{\frac{1}{2}}, \mathcal{G}, \tilde{\mu}), \end{aligned}$$

$$\Pr(\theta_{\tilde{\pi}(t)} = \theta_i, |\mathbf{B} - \mathbf{E}| > \xi^{\frac{1}{2}} \mid \mathcal{G}, \tilde{\mu}) \leq \Pr(|\mathbf{B} - \mathbf{E}| > \xi^{\frac{1}{2}} \mid \mathcal{G}, \tilde{\mu})$$

As $\xi, \epsilon_T \rightarrow 0$, we have shown above that⁵³

$$\Pr(\tilde{\pi}(t) \in \mathcal{T}_L^\xi \mid |\mathbf{B} - \mathbf{E}| \leq \xi^{\frac{1}{2}}, \mathcal{G}), \Pr(|\mathbf{B} - \mathbf{E}| > \xi^{\frac{1}{2}} \mid \mathcal{G}) \rightarrow 0.$$

Further it is straightforward to show that $\Pr(\tilde{\mu}_{\tilde{\pi}(t)} \mid \mathcal{G}) \rightarrow \tilde{\lambda}(\tilde{\mu}(\theta_t(q^t)))$ as $\xi, \epsilon_T \rightarrow 0$, where $\theta_t(q^t)$ is the truthful reporting strategy, and $\tilde{\lambda}(\tilde{\mu})$ is the induced frequency of message $\tilde{\mu}$ when Mediator uses message rule $\tilde{\sigma}(\tilde{\theta}, \mathcal{G})$. Since there are only finitely many messages sent in $\tilde{\sigma}$ (Lemma 1), it is without loss that $\tilde{\lambda}(\tilde{\mu}) > 0$ for any $\tilde{\mu} \in \text{supp}(\tilde{\sigma})$. Thus, using Bayes' rule,

$$\begin{aligned} &\Pr(|\mathbf{B} - \mathbf{E}| > \xi^{\frac{1}{2}} \mid \mathcal{G}, \tilde{\mu}) = \\ &= \frac{\Pr(|\mathbf{B} - \mathbf{E}| > \xi^{\frac{1}{2}} \mid \mathcal{G}) \Pr(\tilde{\mu} \mid |\mathbf{B} - \mathbf{E}| > \xi^{\frac{1}{2}}, \mathcal{G})}{\Pr(\tilde{\mu} \mid \mathcal{G})} \\ &\leq \frac{\Pr(|\mathbf{B} - \mathbf{E}| > \xi^{\frac{1}{2}} \mid \mathcal{G})}{\Pr(\tilde{\mu} \mid \mathcal{G})} \rightarrow 0. \end{aligned}$$

Similarly, $\Pr(\tilde{\pi}(t) \in \mathcal{T}_L^\xi \mid |\mathbf{B} - \mathbf{E}| \leq \xi^{\frac{1}{2}}, \mathcal{G}, \tilde{\mu}) \rightarrow 0$. Thus, as $\xi, \epsilon_T \rightarrow 0$, (32) implies

⁵³Notice that $\Pr(\tilde{\pi}(t) \in \mathcal{T}_L^\xi \mid |\mathbf{B} - \mathbf{E}| \leq \xi^{\frac{1}{2}}, \mathcal{G}) := \mathbb{E} \left[\frac{T_L(\xi)}{T} \mid |\mathbf{B} - \mathbf{E}| \leq \xi^{\frac{1}{2}}, \mathcal{G} \right]$.

that

$$\begin{aligned} & \Pr(\theta_{\tilde{\pi}(t)} = \theta_i \mid \mathcal{G}, \tilde{\mu}) \rightarrow \\ & \Pr(\theta_{\tilde{\pi}(t)} = \theta_i \mid \mathcal{G}, \tilde{\mu}, |\mathbf{B} - \mathbf{E}| \leq \xi^{\frac{1}{2}}, \tilde{\pi}(t) \notin \mathcal{T}_L^\xi) \times \\ & \Pr(|\mathbf{B} - \mathbf{E}| \leq \xi^{\frac{1}{2}}, \tilde{\pi}(t) \notin \mathcal{T}_L^\xi \mid \mathcal{G}, \tilde{\mu}) \end{aligned}$$

But for any time period $\tilde{\pi}(t)$ in which $\Pr(\tilde{\mu}_{\pi(t)} = \tilde{\mu}) > 0$,

$$\Pr(|\mathbf{B} - \mathbf{E}| \leq \xi^{\frac{1}{2}}, \tilde{\pi}(t) \notin \mathcal{T}_L^\xi \mid \mathcal{G}, \tilde{\mu}, \tilde{\pi}(t)) \rightarrow 1$$

since otherwise, $1 - \Pr(|\mathbf{B} - \mathbf{E}| \leq \xi^{\frac{1}{2}}, \tilde{\pi}(t) \notin \mathcal{T}_L^\xi \mid \mathcal{G}, \tilde{\mu})$ would be bounded below by some $\Pr(\tilde{\mu}_{\pi(t)} = \tilde{\mu}) x$, $x > 0$ - a contradiction to

$$\lim_{\xi, \epsilon_T \rightarrow 0} \Pr(|\mathbf{B} - \mathbf{E}| \leq \xi^{\frac{1}{2}}, \tilde{\pi}(t) \notin \mathcal{T}_L^\xi \mid \mathcal{G}, \tilde{\mu}) = 1.$$

Finally, by Bayes' Rule,

$$\begin{aligned} & \Pr(\theta_{\tilde{\pi}(t)} = \theta_i \mid \mathcal{G}, \tilde{\mu}, |\mathbf{B} - \mathbf{E}| \leq \xi^{\frac{1}{2}}, \tilde{\pi}(t) \notin \mathcal{T}_L^\xi) = \\ & \Pr(\theta_{\tilde{\pi}(t)} = \theta_i \mid \mathcal{G}, |\mathbf{B} - \mathbf{E}| \leq \xi^{\frac{1}{2}}, \tilde{\pi}(t) \notin \mathcal{T}_L^\xi) \times \\ & \Pr(\tilde{\mu} \mid \theta_{\tilde{\pi}(t)} = \theta_i, \mathcal{G}, |\mathbf{B} - \mathbf{E}| \leq \xi^{\frac{1}{2}}, \tilde{\pi}(t) \notin \mathcal{T}_L^\xi) \end{aligned}$$

Using an identical argument to that above,

$$\Pr(\theta_{\tilde{\pi}(t)} = \theta_i \mid \mathcal{G}, |\mathbf{B} - \mathbf{E}| \leq \xi^{\frac{1}{2}}, \tilde{\pi}(t) \in \mathcal{T}_L^\xi) \rightarrow \Pr(\theta_{\tilde{\pi}(t)} = \theta_i).$$

Furthermore, for any θ_i , $\xi \rightarrow 0$ implies $\mathbf{b}_i^{\hat{\mathcal{T}}} \rightarrow \mathbf{e}_{\theta_i}$, for $\hat{\mathcal{T}} = \{0, 1, \dots, T-1\} / \mathcal{T}_L^\xi$. By continuity,

$$\lim_{\xi, \epsilon_T \rightarrow 0} \Pr(\tilde{\mu} \mid \theta_{\tilde{\pi}(t)} = \theta_i, \mathcal{G}, |\mathbf{B} - \mathbf{E}| \leq \xi^{\frac{1}{2}}, \tilde{\pi}(t) \notin \mathcal{T}_L^\xi) = \tilde{\sigma}(\tilde{\mu} \mid \theta_i) \quad (33)$$

where $\tilde{\sigma}(\tilde{\mu} \mid \theta_i)$ is the probability Mediator sends $\tilde{\mu}$, given truthful reporting by the sender.

Finally, (33) implies that we can find ξ , ϵ_T small enough that $|\hat{\mu} - \tilde{\mu}| < \eta''$, for any $\eta'' > 0$. Selecting T^* , $\chi(T^*)$ large enough ensures such ξ , ϵ_T can be found. This can be sustained as shown above by a bound on the discount rate, $\delta \geq \delta_{T^*}$ for some $\delta_{T^*} < 1$. \square