Persuasion for the Long-Run*

James Best [†] Daniel Quigley[‡]

December 9, 2016 [Last Updated: 30th January 2017]

Abstract

We examine the limits of persuasion when credibility today is sustained by the incentive of future credibility. We model this as a long-run sender with private information playing a cheap talk game against short-run receivers where there is a noisy signal at the end of each period on the sender's ex-ante private information. We compare our model of long-run persuasion to the persuasion baseline of committed persuasion, where the sender can commit to strategies at the stage game. Long-run persuasion can only achieve the optimal committed persuasion payoffs if the optimal committed persuasion strategy is "honest". When the optimal committed strategy is not honest the use of either a weak communication mechanism called a 'Coin and Cup' (CnC) or a standard communication mechanism (a mediator) expands the Pareto frontier of the game. For sufficiently patient senders, a CnC mechanism replicates committed persuasion payoffs when the sender's information is perfectly observed ex-post, whereas a mediator can get arbitrarily close whenever systematic deviation from truth telling is asymptotically identified. The advantage of the CnC over the mediator is that it is relatively easy to manufacture and implement. Finally, we show how 'emergent communication mechanisms' arise when there are many simultaneous receivers.

1 Introduction

Would be persuaders, such as advertisers, salesmen, media outlets, politicians, central banks, financial analysts, credit rating agencies, lobbyists, think-tanks, charities, activists, employ-

*Special thanks to Bill Zame for his invaluable help on Theorem 1 of this paper. We would also like to thank Peter Eso, Ian Jewitt, Meg Meyer, Adrien Vigier, Joel Shapiro, Ina Taneva and Peyton Young for their comments and advice.

†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)

ment agencies, managers, and many others, rely on their record of honesty to persuade others today. It follows then, that the desire to persuade in the future can generate credibility, and hence persuasion, today.

All this requires that there is some recorded history of honesty, and dishonesty. To this end, many institutions have arisen to keep records of honesty: advertisers fund ombudsmen; central banks have begun publicising the minutes of meetings past; the National Association of Security Dealers lobbied for legislation enforcing the disclosure of historical financial advice; the Dodd-Frank act requires credit ratings agencies to disclose past performance and methodologies; and so on. Further, the internet has created companies that also fulfill this role. For example, Amazon, E-Bay, and Alibaba all use various aggregators of the sellers' claims and buyers' feedback in a manner that increases the persuasive power of some sellers. Another related example, is the current move by social media sites such as Facebook and Twitter to measure and record who produces 'fake news'. Finally, the internet now allows us to use block chain to create decentralised public records of 'claims' and 'feedback' that, in conjunction with modern cryptographic technologies, open up new frontiers in both measuring and recording honesty.

While the improved records generated by such institutions and technologies increase the ability of persuaders to be credible today, it is not obvious that it will make them more honest today. How much then can a persuader gain from being systematically dishonest? How do different institutions and technologies help persuaders use dishonesty more effectively? When are such institutions and technologies socially optimal?

To answer these question we develop a general long-run persuasion game where, contra Kamenica and Gentzkow [2011], there is no commitment at the stage game. A patient Sender ('he') plays a cheap talk game with a sequence of short lived Receivers (each 'she'). Each period a state that is payoff relevant to the Receiver is realised according to an i.i.d. process. The Sender gets a noisy signal of the state; he then sends a message to the Receiver, who then chooses an action based on her belief about the state; finally, there is a public noisy signal about the state. The Receiver's action determines the Sender's payoff. Each new Receiver observes the joint history of messages and public signals which she uses to form her beliefs about the Sender's strategy. The Sender then, must also consider the effect of todays advice on the beliefs and actions of all future Receivers. Hence, the value of being trusted in the future yields a potential source of commitment for the Sender's reporting strategy today.

Our first main result is that, even when the Sender's private information is perfectly observed ex-post, the Sender can only achieve the optimal average payoff if the optimal committed persuasion strategy is "honest". Where a strategy is honest if and only if the Sender's private information determines exactly what message they send, i.e., for each signal they send only one message.[1] The result then follows because dishonest strategies require mixing between multiple messages in some state of the world, and hence indifference across messages. It is helpful to think of the message yielding the lowest stage payoff as "the truth" and the other messages as "lies". Mixing is only then credible if each lie receives an expected punishment that wipes out any gain at the stage from lying. This on path punishment pins the average stage payoff down to the payoff from always sending "the truth" at each stage of the game. Therefore, only when the optimal committed persuasion strategy includes no lies, i.e. only one message for any signal, can long-run persuasion achieve the payoffs available under committed persuasion. By contrast, dishonest equilibria are characterised by trust cycles as Receivers punish the Sender by ignoring him after detecting lies, harming both Sender and Receivers on the equilibrium path.[2]

Next, we look at how we can use different communication mechanisms to increase the efficiency of persuasion. First, we look at a very weak form of communication mechanism that we call a "Coin and Cup" (CnC). A CnC is a payoff irrelevant random variable that the sender privately observes before sending his report each stage. This variable is then publicly revealed only after the Receiver has taken her action. Our second main result is that when there is no noise then the Sender can achieve the same average payoffs as in committed persuasion through use of a CnC. This follows because the Sender never needs to mix with a CnC - instead he can condition his messaging strategy on the value of the CnC at that stage. From the perspective of the Receiver at a given stage the Sender is still mixing messages across states however, yet ex-post everyone can verify that the Sender stuck to his pure CnC conditioned strategy and there is no need for punishment on the equilibrium path.

There are three further things to note about the CnC. First, this device has no effect on the equilibrium of the one-shot cheap talk game. Second, even with noise it still expands the Pareto frontier when optimal committed persuasion is dishonest. Finally, something like the CnC already exists in cryptography - called a "commitment scheme" - that is used for very different purposes. It would be easy to manufacture such a mechanism with a coin and a cup, or a large number of coins and a time-lock safe. However, this is cumbersome, and with the advent of the internet it becomes very easy to embed such a cryptographic device into a block chain that generates a publicly verifiable record of the random variable. It is this simplicity and ease of practical implementation that makes the CnC mechanism appealing.

[1] This notion of honesty allows for coarse messages so long as the message space is a partition of the state space. For example, having a totally uninformative message strategy can be honest i.e. never saying anything. Honesty also includes the standard interpretation of truth telling: Sender's signal can be uniquely determined from his message. Typically the persuasion literature is interested in cases where the optimal committed persuasion strategy is dishonest.

[2] One might tentatively venture that the current episode of voters ignoring experts regarding the Brexit referendum and the US presidential elections is an example of a low trust phase in such an equilibrium.

The third main result is that a standard communication mechanism, or mediator, can get us arbitrarily close to the committed persuasion payoffs so long as the noise satisfies some standard linear independence conditions. The mediator uses a review strategy to switch between persuasion phases and punishment phases. There is a review period with a large known number of stages and a starting point known to the Sender and mediator, but not the receivers. The mediator maps the messages from the Sender into advice for the Receiver. If the Sender reports truthfully, let this mapping yield an expected stage payoff close to the payoff under optimal committed persuasion in the persuasion phase. At the end of the review period the mediator estimates the probability of observing the ex-post noisy signals conditional on the Sender always telling truth. If this probability is low enough the mediator will switch to a punishment phase of babbling for a large finite number of periods. By choosing a long enough review period the probability of punishing a truthful Sender can be made arbitrarily low. The proportion of lies, through this mechanism, can be pushed close to zero. The Receiver's obey the mediator's advice as the advice gives them close to the payoff they would receive under optimal committed persuasion. However, the Sender will still lie, and importantly, they will front load those lies towards the beginning of the review period due to impatience. The mediator's role then, is to aggregate histories so that Receivers don't know where they are in each review period - otherwise, early Receivers don't obey and the equilibrium unravels. A natural example where such mediation strategies are possible would be online sales platforms such as E-bay, where they use bands to classify the quality of sellers - for example Power Sellers - based on the feedback quality over fixed windows of sales.

Finally, we exposit two examples of games which have emergent communication mechanisms. The first is the case of multiple identical receivers acting simultaneously each period with a single state variable. In this game the continuation equilibrium can be conditioned on the actions of all agents in the prior period. The sender can now have a strategy where there is a budget for each message conditional on the Sender's private information, but the sender mixes over who receives the message. This ends up having a similar result to a CnC mechanism. We demonstrate this in a public good game which also shows how the Sender's ability to persuade can be welfare enhancing for all players.

The second case we look at is one where the sender sends simultaneous messages to many receivers acting simultaneously each with their own idiosyncratic payoff relevant state. In this game we can similarly condition the continuation equilibrium on the distribution of actions, allowing a similar budgeting strategy as in the previous case. However, Receiver's have a larger sample to detect deviations from the equilibrium strategy, which reduces the probability of on path punishment occurring due to noise about the Sender's private information. As the number of Receivers becomes large the probability of on path punishment goes to zero. This second case is similar to that of the case of a mediator. We show how this maps into the case of financial analysts and brokerages. In particular, we examine how a legislative change in 2002 pushed by an industry body made it possible to verify the history of advice given by brokerages and how this increased the ability of brokerages to persuade clients.

The rest of this section reviews the literature; section 2 provides a simple example to illustrate the main results of the paper; in section 3 we describe the full model; in section 4 we analyse the equilibria with no noise and the CnC mechanism; in section 5 we analyse the game with noise and the role of a mediator; in section 6 we look at emergent communications mechanisms; and in section 7 we conclude.

1.1 Literature Review

One contribution of our paper is it examines how and when the desire for credibility in a dynamic setting can microfound the commitment assumption presented in Kamenica and Gentzkow [2011] and applied in the subsequent literature (Rayo and Segal, 2010; Kolotilin, 2015; Taneva, 2015; Tamura, 2016). While dynamic persuasion has already been analysed in Rayo and Segal [2010], Kremer et al. [2014], Ely [2015] , Bizzotto et al. [2016] and elsewhere, these analyses all rely on commitment at the stage game. Hence, these papers don't provide foundations for commitment and don't speak to the interaction between credibility now and credibility in the future.

Our paper relates to a recent and growing body of work on repeated Bayesian games with communication and long-run players: Athey and Bagwell [2001, 2008], Escobar and Toikka [2013], Renault et al. [2013], Margaria and Smolin [2015], Hörner et al. [2015]and Barron [2016]. In Hörner, Takahashi, and Vieille [2015], (HTV hereafter), they reduce these games to one-shot games with ex-post transfers - similar to a static Bayesian mechanism design problem - and provide a result analogous to the revelation principle: truth telling does not restrict payoff sets when players are sufficiently patient. However, unlike the cases examined in HTV, non-truthful equilibria expand the payoff set in long-run persuasion games because ex-post transfers are restricted by the impatience of the Receivers. This makes persuasion a meaningful exercise in long-run persuasion games, whereas in these other papers persuasion can be replaced with transfers. Note then, long-run persuasion is to this literature almost as Bayesian persuasion is to mechanism design; a mechanism designer with unconstrained ex-post transfers and no hidden actions can, as in HTV, do away with persuasion.

There are three other papers we know of that drop the commitment assumption of

Bayesian persuasion in one form or another: Chakraborty and Harbaugh [2010], Perez-Richet [2014] and Piermont [2016]. In Chakraborty and Harbaugh [2010] they examine a static cheap talk game, as in Crawford and Sobel [1982], with a multidimensional state and a sender with state independent preferences. They show that the sender can make credible *comparative* statements that trade off the incentives for the sender to exaggerate on multiple dimensions; if sender preferences are quasiconvex they can do better through such persuasion. Unlike our paper they do not microfound commitment, their sender can only make comparative statements, and our results do not rely on the sender having quasi-convex preferences.

In Perez-Richet [2014] and Piermont [2016] the sender observes the probability distribution of states in the future and then chooses an information structure to which they have full commitment. After that the subgame is effectively played by nature with the receiver as nature determines what messages the receiver now gets because the sender has committed to the information structure at all subsequent stages. In these papers then, there is only a lack of commitment in the initial choice of the information structure, once it is chosen they have commitment at every stage of the subgame. The initial choice of the mechanism is pinned down by what that choice reveals about the underlying distribution of the state. This stands in stark contrast to our paper where there is no exogenous ability to commit to the information structure of the game at any point.

2 A Simple Example

The head of a construction Firm wants a Mayor to give them permission for a large construction project. The Mayor is uncertain whether the project will be a net benefit or loss for the city. If the project is rejected both the Firm and the Mayor get a payoff of zero. If the mayor Accepts the project she gets payoff one when the project is Good and minus one when Bad - the firm gets zero either way. Hence, the Mayor will only okay the project if she believes the project will be Good with at least a fifty percent probability. The solid line in Figure 1 below then gives the Firm's payoffs, $v(.)$, as a function of the Mayor's posterior probability, μ, that the project will be Good

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

the Mayor's expected payoff as a function of her posterior belief is given by the diagonal dashed line,

$$u(\mu) = \begin{cases} 2(\mu - 0.5) & if \quad \geq 0.5 \\ 0 & if \quad < 0.5 \end{cases}.$$

The prior probability of the project being Good is $\mu_0 = 1/3$. However, the Firm accurately learns the quality of the project, Good or Bad

Figure 1

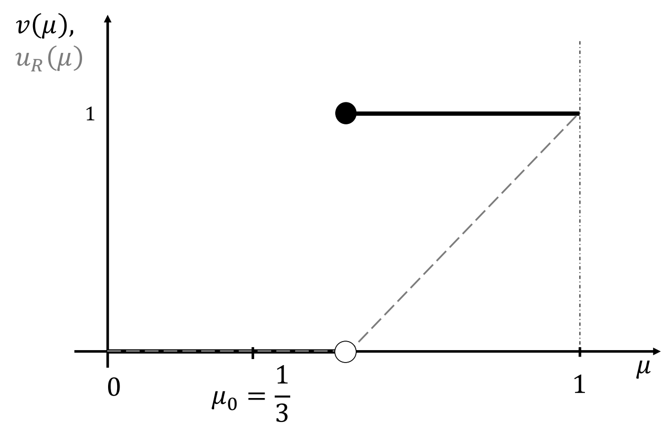

The Firm then sends a report of 'Accept' or 'Reject'. We first consider the committed persuasion case where the firm commits to a policy of sending an 'Accept' or 'Reject' with a project conditional probability, by convention the probability of sending Accept is weakly greater for Good projects than Bad. The (Bayesian) Mayor, after any given report, forms posterior beliefs based on the relative probabilities of receiving that report when the project is Good versus Bad.

The Firm's policy can induce any pair of posteriors that satisfy the law of total probabilities,

$$Pr(Good) = Pr(Good|Accept)Pr(Accept) + Pr(Good|Reject)(1 - Pr(Accept)). \quad (1)$$

Hence, we can think of the Firm's policy as a choice of any two posteriors, $\mu_A = Pr(Good|Accept)$ and $\mu_R = Pr(Good|Reject)$, straddling the prior μ_0; where the posteriors pin down the frequency of sending an Accept. We can restrict attention to policies such that $\mu_A \geq 0.5$ as the Firm can only make profits with policies that give incentive compatible advice where the Mayor accepts the project if she receives an Accept. The expected payoff to the Firm is then just the probability of sending an accept, from (1) this is

$$Pr(Accept) = \frac{\mu_0 - \mu_R}{\mu_A - \mu_R},$$

Figure 2

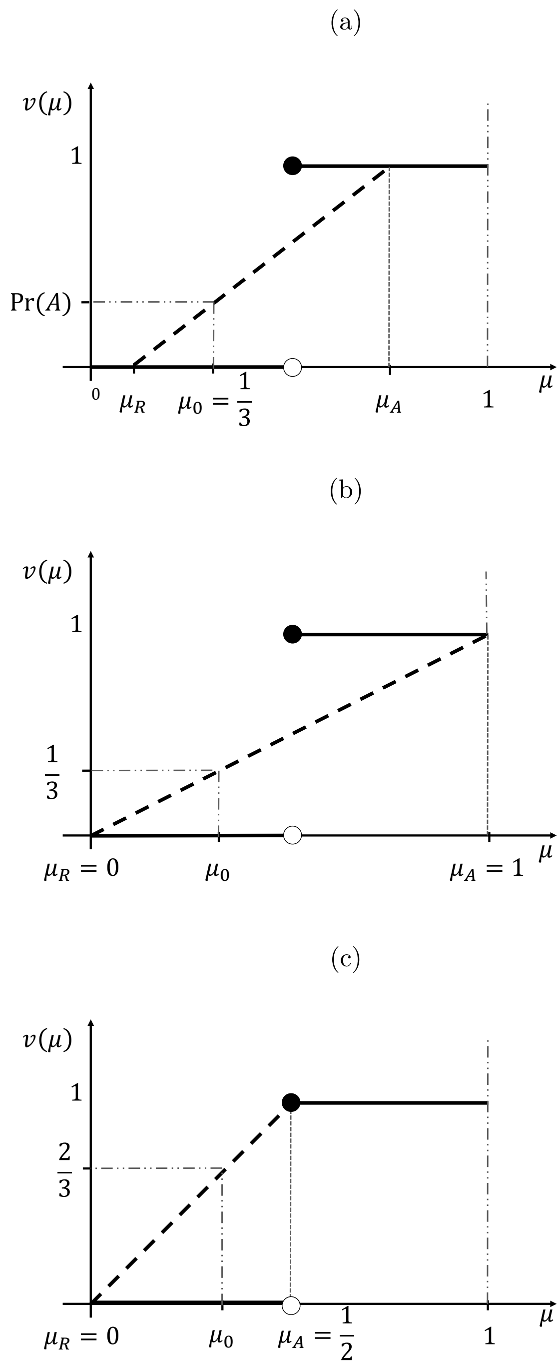

which is just the height of the line connecting the posteriors μ_R and μ_A evaluated at the prior μ_0, as shown in figure 2a below. From the Mayor's perspective, the best policy is Truth Telling - the Firm sends an Accept if and only if the project is Good, $(\mu_R,\mu_A) = (0, 1)$. This pays out for Good projects only - as can be seen from figure 2b, this has an expected payoff of one third to both parties.

From the Firm's perspective this is wasteful, the Firm doesn't need certainty after an Accept to get permission, just $\mu_A \geq 0.5$. Instead the Firm prefers a maximally Persuasive policy where it reports Accept half the time when the project is bad and all the time when the project is Good. Under this policy the Firm is now reporting Accept two thirds of the time and it gives the Mayor posteriors $(\mu_R, \mu_A) = (0, 0.5)$ so the Mayor still accepts after receiving a positive report. While this is the best policy for the Firm it is the worst for the Mayor: the Mayor gets expected surplus of zero irrespective of the report. We can see from figure 2c that this policy places us on the convex hull of the Firm's payoff function, as in Kamenica and Gentzkow (2011).

Commitment at the stage game is frequently infeasible, in which case there is no equilibrium in which the Mayor follows the Firm's advice. The Mayor only accepts if the Firm has a strategy of sending Accept no more than half of the time when the project is Bad. But this can't be an equilibrium as the Firm would always break its word and report Accept. Consequently, without commitment, neither Truth Telling nor Persuasion policies are feasible: all the equilibria of the game are payoff equivalent to a Babbling equilibrium where the Firm randomly sends Accept and Reject reports while the Mayor ignores the reports. However, in a repeated setting, where the Firm cares about future credibility there is room to generate commitment.

In this light, consider now a Firm that is long lived and proposes a sequence of *ex-ante* identical projects to a sequence of one term Mayors. As before, a Mayor at term t accepts the project only if she believes it has at least a fifty percent chance of being Good: her posterior $\mu^t \geq 0.5$. Each Mayor observes all the Firm's prior reports and the outcomes of all prior projects - she forms her beliefs about the Firm's strategy accordingly. The Firm discounts future payoffs at rate δ so that the Firm's lifetime discounted payoff is

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

Figure 3

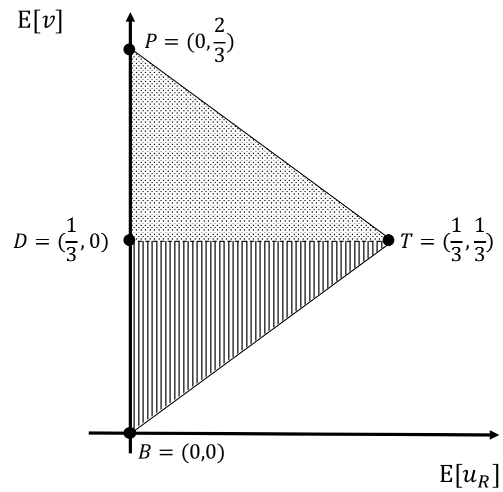

We compare the payoffs of long-term persuasion to those of committed persuasion, described in Figure 3. B is the payoff from Babbling; T is the payoff from Truth Telling; and P is the payoff from optimal committed persuasion. Note, the line PT is the Pareto frontier and B is the worst payoff pair of the stage game.

As Babbling is an equilibrium of the one shot game it is also an equilibrium of the repeated game. The Truth Telling equilibrium can then be supported by the threat of a Babbling equilibrium if the Firm is sufficiently patient. On the equilibrium path, Mayors believe the Firm tells the Truth and accept the project if and only if the Firm sends an Accept report. If the Firm sends an Accept when the project is Bad then the Mayors know they are off path and have the belief that the firm is babbling. The Firm's on path discounted payoff from Truth Telling at any stage is then:

$$V^T = \frac{\mu_0}{1-\delta} = \frac{1}{3(1-\delta)}.$$

If at some term t, the Firm learns the project is Bad but reports an Accept the Firm gains a payoff of one but replaces the Truth Telling continuation equilibrium with Babbling, replacing V^T by zero. Hence, Truth Telling can be supported as an equilibrium if

$$\frac{\delta}{3(1-\delta)} > 1.$$

This result is not surprising. Perhaps more surprising though, is that the Firm can do no better than the Truth Telling equilibrium. Consider an equilibrium where the expected stage payoff is higher than Truth Telling at some stage. Hence, at this stage the Firm's strategy is to sometimes lie when the project is Bad: mix between Accept and Reject so that $\mu_A^t \in [0.5, 1)$. If the Firm is mixing then it must be indifferent between the two reports, as the stage payoff from Accept is higher than Reject it follows that the lower continuation payoff from sending an Accept in the Bad state must exactly offset the higher stage payoff. Suppose then a Firm has a Bad project in some term and it is mixing, but by sheer chance it sends a Reject; then suppose this happens each time it gets a Bad project for which it's strategy is to mix, forever; in this case the Firm never sends an incentive compatible Accept report for a Bad project. The upper bound on the expected payoff conditioned on this accidental outcome of never sending an incentive compatible Accept for a Bad project must then be the Truth Telling equilibrium. Now, the Firm has not lost out by never sending an incentive compatible Accept for a Bad project because the Firm was always indifferent between sending Accept and Reject whenever it was mixing. As the Firm has not lost out it follows that the upper bound on the expected payoff to the Firm is given by the Truth Telling equilibrium.

While the above argument rules out the Firm getting a higher payoff than Truth Telling it does not rule out equilibria in which the Firm lies. For example, consider an equilibrium where the Firm follows the optimal Persuasion strategy in 'normal' periods and babbling in 'punishment' periods. Punishment periods are triggered whenever a Firm lies about a Bad project and go on long enough to make the Firm indifferent between lying and telling the truth. In such an equilibrium the Mayors get zero surplus and the Firm get the same surplus as from Truth Telling. In general, we can any of the average payoffs on the line DT in equilibrium. While it possible to support some degree of persuasion in equilibrium this is Pareto dominated by Truth Telling.

The solution to the Firm's problem is a coin and a cup. At the beginning of each term the Firm 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 Mayor observes the Firm do all this, but does not see the coin. The cup, with the coin still under it, is left on the table. Then, as before, the Firm learns the quality of the project, sends a report, and the Mayor makes her decision. After the decision the Mayor lifts the cup, observes the coin, and records whether it was heads or tails.All Mayors now observe the history of reports, project qualities, and coin flips. This process is repeated every term.

For sufficiently patient Firms, we can replicate optimal committed persuasion. The Firm only ever sends a Reject if the project is Bad *and* the coin comes up tails, otherwise the

Firm sends Accept. On the equilibrium path Mayor t always accepts after the Firm sends Accept, as she has posteriors $\mu_R^t = 0$ and $\mu_A^t = 0.5$. Off the equilibrium path Mayors believe the Firm is babbling. The off path threat of Babbling is enough to stop a patient Firm deviating.

Recall, that without a CnC, mixing between the truth and lying is credible only when there are on path punishments. Now, by introducing the coin and cup, the Firm's pure strategy is still *ex-ante* stochastic for the Mayor making the decision - it is as if the Firm mixes over messages when projects are Bad - yet it is also *ex-post* deterministic.[3] This implies punishments are only required if the Firm deviates from its pure strategy and never on the equilibrium path. Hence, this simple mechanism achieves the maximally persuasive equilibrium with no need for third party verification, contracts, or transfers.

3 The Model

A Sender ('he') and a population of Receivers (each 'she') play the following infinitely repeated persuasion game.

3.1 Stage Game

Each period, a 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, $(\theta_t, \omega_t) \in \Theta \times \Omega = \left\{\theta^1, \theta^2, \ldots, \theta^N\right\} \times \left\{\omega^1, \omega^2, \ldots, \omega^{N_\omega}\right\}$. 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 $\mu_0 \in \Delta\Theta$.[4] Conditional on θ_t, ω_t follows the distribution $g(\omega_t \mid \theta_t) > 0$, where ω_t is also independent of θ_τ, ω_τ, $\tau \neq t$.[5] 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 model 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. Within

[3]The CnC may appear similar to a sunspot, however a sunspot is no substitute as a sunspot contingent strategy would be deterministic *ex-ante* deterministic from the Mayor's perspective: they would see the value of the sunspot and know whether the Firm was going to lie or not for bad projects. The essential feature of the CnC is that the revelation of it's value is staggered, private to the Firm at the point of the Mayor's decision, and only revealed publicly ex-post.

[4]Following standard notation, we use ΔX to denote the simplex over set X.

[5]Note 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 simplifies the exposition of our results related to the review mechanisms we discuss below significantly.

a period, the Sender only cares about the action taken by agent 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 μ_0, $g(\omega_t \mid \theta_t)$. S privately observes the realization of 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 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)* Sender is better informed about Receiver's preferences, $u_R(a,\theta)$;[6] *(ii)* Receiver's final information of her own preferences after taking action will be better than Sender's, $u_R(a,\omega)$;[7] and *(iii)* neither Sender nor Receiver information can be ranked as superior for decision-making, $u_R(a,\theta,\omega)$.[8]

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 $\mathbb{E}[u_R(a_t,\theta_t,\omega_t) \mid m_t]$. Hence, we often refer directly to the Sender's equilibrium period-t stage payoff, as a function of the Receiver's posterior:

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

As in Kamenica & Gentzkow (2011), we assume that whenever R_t's posterior belief leaves her indifferent between two or more actions, we assume 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.[9]

We contrast this stage game to 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:

1a. S chooses an ***experiment***: a message space M, and a random mapping $\hat{s} : \Theta \to \Delta M$.

2a. θ_t is privately drawn from distribution μ_0. Conditional on θ_t, $m_t \in M$ is drawn from s_0.

[6] For instance, where uncertainty is over the quality of a good for sale

[7] This model is most appropriate for experience goods, where consumers have idiosyncratic preferences for goods.

[8] This setting is most appropriate for environments such as asset and/or matching markets.

[9] As we discuss further in Section 4, this is often not the only equilibrium of the stage game.

3a. 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 a stochastic policy. R then observes a draw m_t from the experiment and uses this information to choose an optimal action. Notice in particular that we assume that Sender can only conduct experiments based on his own information, θ_t. This aids comparison with the cheap-talk variant of the game and has the simple interpretation that there is no way for the Sender to predict Receiver's likely feedback, beyond his own private information.

Of course, S can do at least as well using information design as he can in any equilibrium of the static cheap talk game. 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. Kamenica & Gentzkow (2011) show that S's optimal payoff via information design is exactly $\hat{v}(\mu_0)$, which we refer to as *"Optimal Persuasion"*.

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 garbling equilibrium of the cheap talk game. To ensure that persuasion is a useful tool for Sender, we assume in the rest of the paper that v, μ_0 are such that

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

3.2 The Repeated Game

The stage game Γ_t is repeated each period $t = 0, 1, 2, \ldots$, ad infinitum - we refer to this infinitely repeated game by Γ^∞. At each period t and public history $\phi_t = (m_t, a_t, \omega_t)_{\tau=0}^{t-1}$, 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 $a = (a_1, a_2, \ldots)$ 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 $s_t : \Phi_t \times \Theta^t \to \Delta M$ 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 $\rho_t : \Phi_t \times M \to \Delta A$.

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 posterior beliefs $\{\mu_t\}_{t=0}^{\infty}$, 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
$$\mathbb{E}\left[\sum_{\tau=t}^{\infty} \delta^{\tau} u_S\left(a_t\right) \mid \phi_t, \theta_t; \rho\right].$$

2. Given the Sender's strategy, ρ_t maximizes R_t's expected payoff
$$\mathbb{E}\left[u_R\left(a, \theta_t\right) \mid m_t\right] = \sum_{i=1}^{N} \mu_t^i \cdot u_R\left(a, \theta_t^i\right).$$

3. Where possible, the Receiver's posterior beliefs $\mu_t = \left(\mu_t^1, \ldots, \mu_t^N\right)$ satisfy
$$\mu_t^i = \Pr\left(\theta_t = \theta_t^i \mid \phi_t, m_t; s\right)$$

In particular, note that beliefs satisfy the standard "no signalling what you don't know" restriction (Fudenberg & Tirole, (1991)). Regardless of play in rounds $\tau < t$, condition 3. above ensures messages in the support of 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.

3.3 A Direct Equilibrium

Building on an insight from Kamenica and Gentzkow [2011], the notion of equilibrium in our infinitely repeated game can be cast entirely in terms of history-dependent lotteries over beliefs, μ_t. Define the stage game $\hat{\Gamma}_t$ as follows: $\hat{\Gamma}_t$ specifies the Sender's feasible message space as the set of possible posterior beliefs that R_t may hold, $\Delta\Theta$, and is elsewhere the same as Γ_t. The infinitely repeated 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 functions of the form $\sigma = \left(\sigma_\tau\left(h_\tau, \theta^\tau\right)\right)_{\tau=0}^{\infty}$, where each $\sigma_t : H_t \times \Theta^t \to \Delta M$, and $\rho_t : \Phi_t \times M \to \Delta A$ for S, R_t respectively. We denote the set of all strategies for S by Σ.

We define a *direct equilibrium* of this repeated game as follows:

1. *(Best responses)* Given the Receivers' belief functions $\mu_t(h_t, \tilde{\mu}_t)$, $\tilde{\mu}_t \in \cup_{\theta_t \in \Theta} supp(\sigma_t(h_t, \theta_t))$ maximizes the Sender's expected discounted payoff

$$V_t(h_t, \theta_t) = v(\mu_t(h_t, \tilde{\mu}_t)) + \delta \mathbb{E}[V_{t+1}((h_t, \tilde{\mu}_t, \theta_t), \theta_{t+1})] \tag{2}$$

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

2. (*Obedient beliefs)* The Receiver believes any equilibrium message, $\tilde{\mu}_t \in \cup_{\theta_t \in \Theta} supp(\sigma_t(h_t, \theta_t))$

$$\mu_t(h_t, \tilde{\mu}_t) = \tilde{\mu}_t$$

3. *(Bayes plausibility)* $\mu_0 \in co(\cup_{\theta_t \in \Theta} supp(\sigma_t(h_t, \theta_t)))$.

The function $\mu_t(h_t, \tilde{\mu}_t)$ specifies R_t's beliefs, given observation of history h_t and message $\tilde{\mu}_t$ sent by S in period t. Given these beliefs, the optimal behaviour of the Receiver is implicit in the function $v(\mu_t(h_t, \tilde{\mu}_t))$, which defines S's stage payoff from this behaviour. V_t is simply the sum of S's discounted payoff from equilibrium play, from history (h_t, θ_t) onwards. In any equilibrium, S must maximize (2) 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' the the Receiver's prior. While these two conditions appear stronger than required for any equilibrium, the following Lemma establishes that it is without loss to restrict attention to such *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.*

Lemma 1 extends the insight of Kamenica & Gentzkow (2011) to equilibria of repeated games, in which Sender is unable to commit to his signalling strategy at any history.As for condition 2, since all that matters about S's strategy is the effect it has on R_t's beliefs, we can replace S's messages with recommendations of the beliefs that R_t should hold. Moreover these recommendations must be optimal for S, since they were in game Γ^∞ (condition 1).[10] Finally, condition 3 follows from the observation that any equilibrium strategy must induce posteriors that satisfy the Law of Total Probability, (**??**). We sometimes write an induced lottery of posteriors, at history h_t, as $\lambda \in \Lambda(\mu_0) \subset \Delta(\Delta\Theta)$, where $\Lambda(\mu_0)$ denotes the set of lotteries satisfying condition 3. In other words, if at some history h_t, λ induces a lottery

[10] For messages not in the support of a continuation equilibrium at h_t, R_t's beliefs are not constrained by Bayes' rule and thus it is easy to prevent Sender from playing 'off path' choices of $\tilde{\mu}$.

of M posteriors $(\mu_{t,1}, \mu_{t,2}, \ldots, \mu_{t,M})$ with $\lambda_j = \Pr(\mu = \mu_{1,j})$, $j = 1, 2, \ldots, M$ and $\lambda \in \Lambda(\mu_0)$, then it satisfies

$$\mu_0 = \sum_{j=1}^{M} \lambda_j \mu_{t,j} \tag{3}$$

4 The Value of Repetition for Persuasion

In this Section, we are primarily interested in understanding when the opportunity for repeated interaction can allow S to achieve his optimal discounted average payoffs under *persuasion*, despite only being able to make cheap talk statements. First, we establish that there is generally a need for repeated interaction to improve the possible payoffs that Sender can achieve in equilibrium. Interestingly, it is possible to find preferences $v(\mu)$ for which optimal persuasion can be achieved as an equilibrium of a static cheap talk game with non-trivial communication (see Figure 4). However, our first main result establishes that these kind of functions are not typical and therefore static cheap talk does not usually allow Sender to do as well as he would under commitment to optimal persuasion:

Theorem 1. *For any prior μ_0, optimal persuasion is generically not possible in a static cheap-talk game.* [11]

While there exist non-convex functions $v(\mu)$ for which the concavification $\hat{v}(\mu)$ involves at least two points μ_x, μ_y such that $v(\mu_x) = (\mu_y)$ and $\hat{v}(\mu_{x,y})$, $\forall \mu_{x,y} = \alpha\mu_x + (1-\alpha)\mu_y$, $\alpha \in [0,1]$. In such cases, S is indifferent between sending messages μ_x and μ_y and moreover these messages are feasible in a direct equilibrium if $\mu_0 = \mu_{x,y}$ for some $\alpha \in [0,1]$. Thus, there exists an equilibrium of the static cheap talk game which achieves $\hat{v}(\mu)$ (and therefore, S's optimal stage payoff under commitment[12]) without the need for repeated play. The proof of Theorem 1 shows that such functions are in fact non-generic. Since such cases are rare, we focus in the rest of the paper on functions v and priors μ_0 for which no cheap talk game can achieve S's optimal payoffs under commitment to persuasion.

Theorem 1 tells us that cases in which Sender can achieve optimal persuasion using cheap talk without the need for repeated interaction are rare. Moreover we are interested studying the role of repetition in persuasion. We therefore focus our attention on the generic cases of Sender payoff function for which static cheap talk cannot be used to sustain optimal persuasion.

For any subset $P \subseteq \Theta$, define $\Delta_P\Theta := \{\mu \in \Delta\Theta : \theta^i \in P \iff \mu^i = 0\}$ as the set of posteriors which put positive probability on a state if and only if it is in P.

[11] Special thanks to Bill Zame for advice on this proof.

[12] This follows immediately from Corollary 2, Kamenica & Gentzkow (2011).

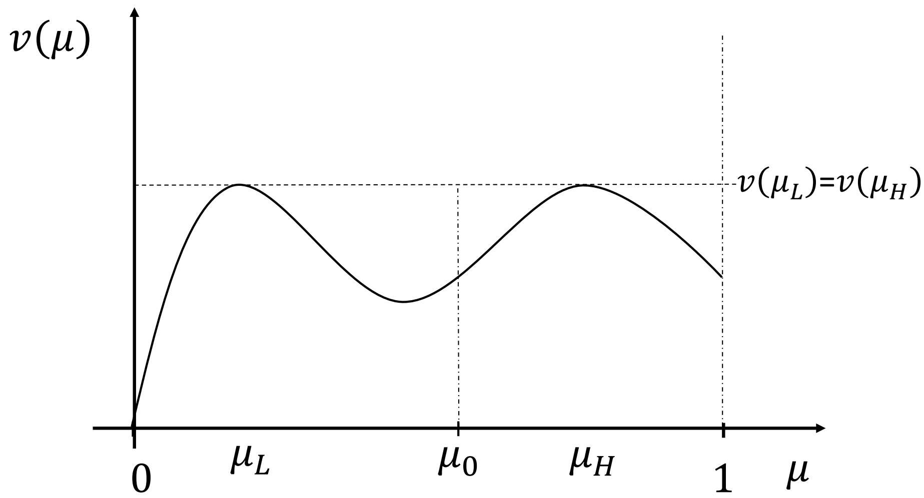

Figure 4: A (non-generic) example of optimal persuasion attainable via cheap talk

Assumption 1. *For all $P \subseteq \Theta$, optimal persuasion cannot be achieved by informative, static cheap talk, conditional on $\theta \in P$.*

In the Appendix, make the formal statement of Assumption 1. Assumption 1 ensures that there is no subset of types for which S could use one-shot, informative communication to achieve the optimal commitment payoff $\hat{v}_P$, conditional on R also knowing that $\theta \in P$. Considering first $P = \Theta$, Theorem 1 assures us that functions violating Assumption 1 in this case are non-generic. In other words, such functions are rare. Since Theorem 1 applies to any finite state space, Θ, we can similarly apply the logic to the payoff function v, defined over the subspace $\Delta_P\Theta$. In this way, Theorem 1 also assures us that v functions for which cheap talk could achieve the optimal commitment outcome for S on any subset P of the state space are also rare. Since these cases are non-generic, we omit their analysis from the main results in order to aid exposition of the typical persuasion problem.

By ruling out only *informative* cheap talk as a method of achieving the KG (2011) solution on any $P \subset \Theta$, Assumption 1 allows for situations in which S prefers to communicate no more information if he knows that R believes $\theta \in P$. This is important, as we do not wish to rule out cases in which S never wishes to conceal information - indeed, strategically concealing information is at the heart of persuasion. Our assumption allows for this. In fact, as we have emphasized it allows for almost any combination of interior and boundary beliefs as part of the optimal commitment signal. It only requires that such signals do not all leave S indifferent when induced in the stage game.

Before we analyse the properties of the repeated game in detail, we present a useful preliminary result:

Lemma 2. *In any equilibrium, Sender can do no better than a strategy which at any history h_t induces at most N possible posterior beliefs, $\mu_t(h_t)$.*

Lemma 2 establishes that from the perspective of S's payoffs it is without loss to restrict attention to direct equilibria in which at any history, S's strategy induces no more than an N-point distribution over posterior beliefs (recall that $N = |\Theta|$). The result simplifies the search for optimal equilibria significantly. Most importantly, it ensures us that we only need consider strategies that induce a finite distribution of posteriors for any Receiver, R_t. We use some key properties of direct equilibria for S's payoffs and of convex sets to establish that if $N' > N$ signals were ever being sent at some history h_t, one of these signals would be redundant for S's continuation payoff at that history (and for feasibility of induced posteriors at that history). Removing such an alternative from S's strategy at h_t is feasible at h_t, and since it does not affect payoffs at h_t, it does not affect S's incentives at earlier histories or indeed his expected discounted payoff from the game, $\mathbb{E}_\theta[V_0(\theta_0)]$. Interestingly, the properties of equilibrium allow us to reduce the cardinality of the signal space by more than under standard persuasion, which can reduce the search over signals to $N+1$-point distributions.[13]

4.1 Repeated Persuasion without Noise

In this section, we provide a first analysis for the case where Sender's private information is a perfect prediction Receiver's finally preferences: that is, we assume $\Omega = \Theta$ and $g(\theta_t \mid \theta_t) = 1$. Intuitively, because Sender's type can be verified at the end of each round, this environment represents the very best case scenario for the use of reputation as a device for allowing Senders to convincingly persuade a population of Receivers. Somewhat surprisingly, however, we show that repeated interaction cannot typically achieve the optimal commitment payoff from persuasion *even* when there is no noise ex post between the information of Senders and Receivers. We extend our results to allow for Sender's information to be only a noisy signal of Receiver's final utility in the next section. Suppose now S has the opportunity to interact sequentially with a (potentially infinite) set of short-run Receivers. As a preliminary result, we show that repeated play of the cheap-talk game can sustain truth-telling by the Sender as an equilibrium.

[13] The proof deals with two complications as compared with standard persuasion arguments. First, it deals with the fact that S does not commit to his strategy. And second, we must ensure that the equilibrium dynamics are not violated.

Proposition 1. *There exists $\overline{\delta} < 1$ such that truth-telling is an equilibrium of the repeated game $\forall \overline{\delta} \leq \delta < 1$, iff Sender's truth-telling payoff exceeds his worst stage game equilibrium payoff.*

To sustain on-path truth-telling in every period the equilibrium employs a trigger strategy, moving to the worst cheap-talk equilibrium forever if a deviation is detected[14]. When θ_t can be observed at the end of each round, deviations from truth-telling are easily detectable to Receivers. Therefore, so long as Sender is sufficiently patient and the worst cheap talk equilibrium yields a lower expected Sender stage payoff than does truth-telling, such a strategy enforces truthful equilibria.

We now ask how the potential for using repetition as a commitment device affects S's ability to earn rents from persuasion in the equilibrium of some repeated cheap talk game. In asking this question, we move to a focus on Sender-preferred equilibria.

In particular, consider the problem of maximizing S's period-0 discounted utility, across all possible equilibria of the repeated cheap talk game:

$$\max_{\sigma \in \Sigma} \mathbb{E}_{\theta} \left[V_0 \left(\theta_0 \right) \right] \tag{4}$$

s.t.

$$V_t \left(h_t, \theta_t \right) \;\; = \;\; v \left(\mu_t \right) + \delta \mathbb{E} \left[V_{t+1} \left(\left(h_t, \mu_t, \theta_t \right), \theta_{t+1} \right) \right] \geq v \left(\mu_t' \right) + \delta \mathbb{E} \left[V_{t+1} \left(\left(h_t, \mu_t', \theta_t \right), \theta_{t+1} \right) \right],$$

$\forall h_t \in H_t, \mu_t \in supp(\sigma_t \left(h_t, \theta_t \right))$, $\mu_t' \in \cup_{\theta_t \in \Theta} supp(\sigma_t \left(h_t, \theta_t \right))$, and

$$\mu_0 \in co \left(\cup_{\theta_t \in \Theta} supp \left(\sigma_t \left(h_t, \theta_t \right) \right) \right),$$

$\forall h_t \in H_t$.

Problem (4) involves choosing a strategy profile $\sigma = \left(\sigma_1, \sigma_2 \left(h_2 \right), \dots \right)$ for S that maximizes his present discounted utility, subject to: *(i)* each choice of $\mu_t \in supp(\sigma_t \left(h_t, \theta_t \right))$ involves a (weakly) higher present discount value for S at history h_t than any alternative μ_t' that is played by S with some positive probability at h_t; *(ii)* satisfying Bayes plausibility at each history. 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 believing that these choices be consistent with the equilibrium strategy, σ. In particular, non-equilibrium choices of $\tilde{\mu}$ can be ruled out because we can choose R_t's beliefs

[14]Such an equilibrium can in general be worse for Sender than babbling.

to be skeptical after such reports.[15] 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.

Next, we characterize the solution to the value function for problem (4). First, we introduce some notation. Let $\underline{v}_i(\lambda) := \min\{v(\mu) : \mu \in supp(\lambda), \mu^i > 0\}$ be the minimum payoff to S among all posteriors μ that *(i)* are in the support of N'-point distribution $\lambda \in \Delta(\Delta\Theta)$, for $N' \leq N$, and *(ii)* occur with strictly positive probability conditional on state θ_t^i (under λ). Then we have:

Proposition 2. *Sender's discounted average continuation value from any repeated cheap talk game is bounded above by*

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

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

Proposition 2 establishes an upper bound on the payoffs that Sender can achieve in any equilibrium of the repeated game. Moreover, it shows that this upper bound is attainable in some equilibrium, so long as Sender is sufficiently patient. The bound in equation (5) states that Sender's best discounted average payoff must be no greater than the best expected statewise-minimal payof, among all lotteries of posteriors $\lambda \in \Lambda(\mu_0)$. The importance of the statewise-minimal payoff is as follows: in any equilibrium, if there are multiple messages in the support of Sender's strategy conditional on some θ, then Sender must be held indifferent across these messages. Thus, messages which lead to more preferable current actions must also be associated with larger future punishments. In a Sender-preferred equilibrium, the least-preferred current action is never associated with future punishment, pinning down the upper bound on Sender's equilibrium discounted average payoffs. In the Appendix, we construct a strategy profile which achieves this bound using finite punishment periods for sending messages other than the Sender's least-preferred one at some θ_t, followed by reversion to the original strategy thereafter. Since the worst cheap talk equilibrium yields statewise minimal payoffs, it is necessarily (weakly) worse than the bound in (5), and thus may always be used as a punishment (or itself achieves the Sender's largest payoff).

In comparing the repeated cheap talk game to the commitment benchmark, a particular class of experiments is important for understanding when the bound on S's payoff, (5),

[15] To support the equilibrium, we additionally specify the continuation play after sending $\tilde{\mu}$ as equal to the worst on-path message at h_t, thereafter.

is tight, relative to the commitment payoff $\hat{v}(\mu_0)$. We introduce the notion of an *honest* experiment, as follows:

Definition 1. An experiment $(\hat{s}, M)$ (resp. behavioural strategy, $(\hat{\sigma_t}, M)$) is *honest* if there exists a partition P of Θ such that $\hat{s}$ (resp. $\hat{\sigma_t}$) can be expressed as a bijection $\hat{s}(\hat{\sigma_t}) : P \to M$.

Under an honest reporting strategy there is no $\theta \in \Theta$ for which S might randomize over sending two or more messages. Honest experiments (reporting strategies) are therefore ones for which each message convinces the Receiver that θ lies in a different, disjoint subset of Θ. In other words, assuming S plays according to $\hat{\sigma}_t$ at round t, R_t receives a truthful report about the element of the partition P in which θ lives. While these reports are truthful about the partition in which θ lies, they are nonetheless coarse signals. In particular, our definition of honesty allows for completely uninformative 'babbling' experiments (reports) and perfect truth-telling.

With this definition in hand, we can establish the following Theorem:

Theorem 2. *Suppose Assumption 1 holds. Optimal persuasion is attainable by repeated cheap-talk if and only if the optimal experiment is honest, for* $\delta \geq \underline{\delta}$.

Theorem 2 tells us that S's optimal persuasion payoff $\hat{v}(\mu_0)$ can be attained in the equilibrium of the repeated cheap-talk game if and only if the optimal experiment under commitment involves sending messages which are *honest*. The intuition for this result rests on a simple observation: Receivers can verify ex post whether S has deviated from making reports that are consistent with honest experiments. Therefore, honest reporting strategies can be sustained in equilibrium of the repeated game by using only the *threat of off-path punishments*. If the optimal experiment under commitment happens to be honest, then there is an equilibrium of the game in which S makes reports that mimic this experiment on-path, sustained by off-path punishments whose costs are never realized (for large enough δ) . By contrast if the optimal commitment experiment only involves randomization between messages at some θ, S's strategy can never mimic such an experiment at any history without leaving Receivers in doubt about whether S has deviated from the experiment at some θ_t. In order to ensure incentive compatibility, Receivers thus need to punish S on the equilibrium path following some messages.

4.1.1 Recovering Sender Optimality: 'Coin and Cup' Mechanisms

As we have seen in Section 4, the repeated opportunity for cheap talk does not typically allow Sender to achieve his optimal commitment payoff $\hat{v}(\mu_0)$. However, Theorem 2 also provides hope that there might be ways to design institutions such that S can achieve $\hat{v}(\mu_0)$

in some equilibrium. As we pointed out above, a key feature of *honest* experiments is that in each period Receivers can verify ex post whether S has deviated from making reports that are consistent with such experiments. Importantly, this allowed for on-path equilibrium strategies to be sustained using only off-path punishments. Thus, if we can find mechanisms which allow for this feature without also insisting on strict honesty (with respect to Θ) then we might be able to recover optimal long-run equilibria in these games too.

In this section, we introduce a simple 'Coin and Cup' (CnC) mechanism which can indeed be used to ensure S can achieve $\hat{v}(\mu_0)$, without violating incentive compatibility. The CnC mechanism augments the repeated game in a payoff irrelevant way for all players but nonetheless introduces equilibria that attain $\hat{v}(\mu_0)$ for S. In addition these mechanisms are robust to changes in the payoffs of Sender and Receiver. Finally we identify several examples of real-world applications in which institutions appear to use naturally occurring versions of CnC mechanisms and show how these institutions can be used to maximize Sender's payoffs.

A 'Coin and Cup' mechanism introduces a payoff-irrelevant state variable, $c_t \in [0,1]$, to the repeated game in Section 3. To ease exposition, we suppose that c_t is i.i.d. over time $c_t \sim U[0,1]$. The CnC mechanism is a combination of sequence of random variables c_t and repeated play of the following adapted stage game:

1a. θ_t, c_t are drawn independently from their respective distributions. S privately observes ***both*** realizations.

2a. S sends a message $m_t \in M_t$ to R_t

3a. After observation of m_t, R_t chooses an action $a_t \in A_t$.

4a. After taking action a_t, c_t and the state θ_t are observed by all Receivers.

The CnC mechanism requires that we can find a payoff-irrelevant random variable each period such that *(i)* Sender is able to privately observe c_t before R_t; *(ii)* Sender cannot manipulate the realization of c_t; *(iii)* R_t is able to observe c_t. In this way, one can think of c_t as a staggered sunspot.

As the next Proposition shows, the CnC mechanism always admits equilibria in which Sender achieves his full commitment payoff:

Theorem 3. *Suppose Sender can achieve payoff $\nu^\star$ via commitment to some experiment. $\nu^\star$ is attainable in an equilibrium of some CnC mechanism, if $\nu^\star$ exceeds S's worst stage game payoff.*

Intuitively, the CnC mechanism improves on pure repetition because it allows punishments to be conditioned on the realization of a much larger 'augmented' state, (θ_t, c_t), where c_t does not involve any new payoff considerations for Sender or Receiver. On this expanded state space, we can use c_t as a way to assign a 'budget' for reports sent given each θ_t. Thus, if the optimal experiment ever proscribes mixing between two messages μ_1, μ_2 with probabilities λ_1, λ_2 at some state θ'_t, we can simply choose the equilibrium strategies to allow message μ_1 to be reported without subsequent punishment if $c_t \in [a, a+\lambda_1]$, $a+\lambda_1 \leq 1$, and μ_2 to be reported if $c_t \in [b, b+\lambda_1]$, $b+\lambda_1 \leq 1$, for disjoint intervals $[a, a+\lambda_1]$, $[b, b+\lambda_1]$.[16] Importantly, this means that each round Sender can credibly mimic the optimal experiment, knowing that if he reports according to his 'budget' (given θ'_t), he will face no punishment. Otherwise, he will face a 'Grim Trigger'-style punishment in which future Receivers all revert to the babbling equilibrium.

Importantly, for the CnC mechanism to improve equilibrium outcomes, Sender must know the realization of c_t *before* reporting (to know which message to send without punishment), while Receiver must only observe c_t *after* taking action a_t. If Receiver learned c_t too early, this would destroy Sender's ability to effectively persuade without facing on-path punishments. In other words, the environment would revert to the repeated games setting of Section 3. Thus, the ability to delay observation of c_t to Receivers is a crucial element of the design of CnC mechanisms.

Such staggered sunspots c_t are implementable and don't require any specialist knowledge of the decision problem or for instance the realization of θ.[17] As the simplest possible example, we can create them using no more than some coins and a cup. So long as the Receiver observes the coins being tossed into the cup, Sender can privately peer in and check if each is a 'Heads' or a 'Tails', before sending his messages to Receiver.[18] The proportion of 'Heads' across the cups can then play the role of c_t. As a more realistic example, Blockchain technologies (such as that underpinning Bitcoin) support decentralized recording and updating of information among peers using cryptographic methods.[19] 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.[20] Programmed

[16]Since mixing probabilities sum to 1 and ω_t is uniform, it is easy to characterize disjoint sets intervals on Ω that support such a strategy. A similar logic goes through for more general distributions of ω, so long as the distribution at each t is atomless.

[17]Notice that we do not require S to be able to commit to a specific experimental procedure for generating a particular distribution, ω_t.

[18]Sender is not allowed to touch the cup.

[19]Other recent uses of Blockchain technology include: 'smart contracts' for verifying the performance of obligations, reducing manipulation of experimental design in medical trials, and creating trustworthy digital accounts of property ownership.

[20]http://www.economist.com/news/science-and-technology/21699099-blockchain-technology-could-

with a random number generator to update the Blockchain with new messages c_t, these technologies could be used as the basis of a CnC mechanism. Interestingly, the availability of 'trust' technologies such as Blockchain can therefore introduce equilibria in which the payoffs to persuasion are improved.

5 Repeated Persuasion with Noise

We now return to the more general setting of Section 3, in which Receivers can never verify the information on which Sender based his recommendations. In this setting, it is now typically not possible to verify ex post whether a Sender has played according to a proscribed strategy, even if the optimal signal under commitment is honest (or if the coin and cup technology is available. Moreover, given a realized θ_t, it is usually not even possible to determine whether Sender has made the 'worst' report in the conditional support of $\sigma_t(h_t, \theta_t)$.[21] To ensure the Sender is indifferent between messages in the support of his equilibrium strategy therefore requires even greater punishments on the equilibrium path. In particular, at time $t$punishments must now be assigned in expectation across different realizations of ω. As a result, the 'worst' message associated with some state (e.g. 'Accept' when $\theta_t = Good$ in the Mayor-Firm equilibrium, P) will typically have to sometimes be proscribed an 'accidental' punishment, meant for that message in another state (e.g. 'Accept' when $\theta_t = Bad$ in equilibrium, P). Accordingly, not only does equation (5) continues to provide an upper bound on payoffs but the bound is no longer always achievable in the repeated game with noise. To the extent that punishment periods may also be costly for Receivers (as is true when reversion to periods of babbling acts as the punishment) the costs associated with 'accidental' punishments can also fall on Receivers.[22] This discussion is summarized in the following remark:

Remark 1. In the presence of noise, Sender's payoff from repeated persuasion is no better, and can be strictly worse, than without noise.

The key problem introduced by idiosyncrasies in Sender and Receivers' types is that ex post verification is no longer possible. As a result, some punishment is required to discipline Senders from 'over-sending' messages that lead to attractive actions. In this environment, where the costs of sustaining punishments are large, it is arguably even more important

improve-reliability-medical-trials-better

[21]Except where Sender's recommendation induces his least favoured action among *all* possible messages in $\cup_{i\in\{1,\ldots,N\}} supp(\sigma_t(h_t, \theta_t^i))$.

[22]However, this is not guaranteed to be costly for Receivers - as discussed above, there are typically multiple Nash equilibria of the stage game, some of which may offer Receivers a higher payoff than they achieve 'on the path'.

to find communication devices like the ‘Coin and Cup’ that can restore credible persuasive equilibria without the need for punishment phases. As our next remark points out, while the CnC mechanism cannot typically completely overcome the problem of ex post unverifiable reporting and recover the full commitment payoff for the Sender, it does still widen the set of available payoff pairs available to Sender and Receivers.

Remark 2. The CnC mechanism allows us to *strictly* expand the Pareto frontier whenever an equilibrium inducing at least one interior posterior $\mu_t \in int\,(\Delta\Theta)$ lies on the Pareto frontier of the repeated game.

Of course, the CnC can never make the set of equilibrium payoffs smaller for Sender and Receivers, since it can always be ignored. As an illustrative example, consider a variant of the Mayor-Firm game in which $\Pr(\omega = good \mid Good) = \Pr(\omega = bad \mid Bad) = p > 0.5$. Suppose that in the Sender-optimal equilibrium, the Firm’s equilibrium strategy proscribes reporting ‘Accept’ with probability 1 if $\theta_t = Good$, and ‘Accept’ with probability $q \leq 0.5$ if $\theta_t = Bad$. To keep Sender indifferent between reports when $\theta_t = Bad$, such a strategy requires ‘accidental’ punishments (following $\omega = bad$) even when Sender says ‘Accept’ in the Good state. The Coin and Cup can help reduce punishments here, without affecting reporting incentives. The equilibrium simply allows Sender to report ‘Accept’ without punishment (regardless of ω) for $0 \leq c_t \leq q$. Otherwise, continuation play is as before. Since Sender privately observes c_t upfront, this change does not affect his incentives for $c_t > q$ and he can report ‘Accept’ if and only if $\theta_t = Good$. However, for $0 \leq c_t \leq q$, he always prefers to say ‘Accept’ (therefore matching the overall distribution of reports from the repeated game) in which case he avoids punishment regardless of θ_t.

However, in general the CnC mechanism will not allow us to retrieve the full set of payoffs available under commitment. This raises the question of whether a trusted mediator can improve the scope of communication possibilities between the Sender and Receivers by carefully controlling the information that Receivers observe about the Sender’s reputation. Such mediated trust mechanisms are common in the real world, most notably in the variety of online platforms such as eBay, Alibaba, Amazon, AirBnB and others. On these websites, independent sellers are given ratings by the platform based on the noisy aggregation of buyers’ feedback about several aspects of their experience with the seller, including the quality of the good or service purchased, the accuracy of the seller’s claims etc.

In order to address this question, we now describe a particular mediated communication mechanism, which we call a Review Mechanism, since it is based on Radner [1985]’s review strategies. A formal description of the mechanism is postponed to Appendix B and here we settle for an intuitive description. In the Review Mechanism, the Mediator controls the flow of information available to Receivers. The mechanism runs iteratively in phases. At time 0, the mechanism starts in a 'Good' phase, which lasts for T periods. At each period within this phase, a receiver is only told that the Sender's past play means that his reputation is 'Good', $\mathcal{G}$. Importantly, Receivers observe *nothing else*. In particular, we assume that the Mediator can randomly permute the ordering of the first T Receivers so that they do not even know their position in the line, and cannot infer this from their index t. During phase $\mathcal{G}$, the Mediator privately asks Sender for reports of his type, and commits to garble these reports using a message function $\sigma_m(\theta, \mathcal{G}) : \Theta \to \Delta\Theta$ which outputs recommended beliefs, $\tilde{\mu}$, for the Receiver who is assigned to play at time t. Sender observes nothing else about the mechanism other than the phase $\mathcal{G}$ and his history of reports and induced actions. In particular, he does not observe $(\omega_\tau)_{\tau=0}^{t}$. At the end of the first T periods, the mediator conducts a statistical review of Sender's reports, $\left(\tilde{\theta}_t\right)_{t=0}^{T-1}$, given the sequence of feedback $(\omega_t)_{t=0}^{T-1}$. The test checks whether the Sender is sufficiently likely to have reported truthfully on average across the $\mathcal{G}$-phase. If he passes, the mechanism moves onto another $\mathcal{G}$-phase from period T onwards. Otherwise, mediator switches to announcing to the next $\beta T \in \mathbb{N}$ Receivers that the the Sender's reputation is $\mathcal{B}$. In this case, the Mediator commits to send a 'babbling' message to each Receiver, regardless of Sender's private reports in the $\mathcal{B}$-phase. At the end of the $\mathcal{B}$-phase, mediator reverts to announcing a $\mathcal{G}$-phase.

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 to the joint probability of $\left(\theta_t^i, \omega_t^j\right)$, $i, j \in \{1, 2, \ldots, N\}$ by $f\left(\theta^i, \omega^j\right) := \mu_0^i g\left(\omega_t^j \mid \theta_t^i\right)$. Then:

Assumption 2. *There exists a subset $\Omega' := \left\{\omega^1, \ldots, \omega^N\right\} \subset \Omega$ such that*

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

are linearly independent.

Assumption 2 requires that while the Receiver cannot typically infer Sender's signal θ_t from observing her own ω_t, each possible signal provides statistically distinct posterior distributions $\theta \mid \omega^i$ in the sense that such a posterior could be formed only by ω^i and not some linear combination of signals, $\{\omega^j\}_{j \neq i}$. To understand this condition better, consider the following example:

Example 1. *(Symmetric Binary Signals)* Suppose Sender's signal at time t is binary, $\theta_t \in \left\{\underline{\theta}, \overline{\theta}\right\}$, distributed according to $\Pr\left(\theta_t = \overline{\theta}\right) = \mu_0 \in (0,1)$. Receiver R_t observes $\omega_t \in \{\underline{\omega}, \overline{\omega}\}$, where $\Pr\left(\omega_t = \overline{\omega} \mid \overline{\theta}_t\right) = \Pr\left(\omega_t = \underline{\omega} \mid \underline{\theta}_t\right) = p$. It is easy to verify that Assumption 2 is satisfied iff $p \neq \frac{1}{2}$. In other words, Assumption 2 is satisfied so long as ω_t is an informative signal of θ_t, in that $\Pr\left(\overline{\theta}_t \mid \omega_t\right) \neq \mu_0$.

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

Assumption 3. *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 $\lambda' \in cl(X(\lambda'))$.*

Assumption 3 ensures that the Sender's reduced-form utility function does not involve discontinuous jumps in Sender's payoffs which are only attainable on a lower-dimensional subset of beliefs in $\Delta\Theta$. Since u_S is continuous in actions in our model, such 'jumps' can only ever occur if the Receiver's beliefs just induce him to change his actions in a discontinuous way, as his beliefs enter such 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, such utility functions 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.

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 Sender's private information is not verifiable ex post:

Theorem 4. *Suppose Assumptions 2 and 3 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 Sender's normalized discounted average utility is at least $\hat{v}(\mu_0) - \varepsilon$:*

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

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

Theorem 4 establishes that trust mechanisms, such as the mediated review systems on sales platforms like eBay, are a valuable tool for persuasion purposes. They can allow Senders to extract almost all the available gains from persuasion even when the Sender cannot commit to signals, and where repeated play alone bounds Sender's payoff by (5), which is often strictly below $\hat{v}(\mu_0)$.

To ensure an equilibrium in which the Sender achieves 'close to' optimal payoffs, the proof constructs a Review Mechanism with the features that *(i)* the mediator promises to use a garbling function $\sigma_m(\cdot, \mathcal{G})$ 'close to' the optimal experiment under commitment, where recommendations $\tilde{\mu}_i$ are each at continuity points of v (guaranteed by Assumption 3); *(ii)* the test is strict enough that Sender's optimal strategy can be bounded 'close enough' to truth telling by the expected costs of falling into a $\mathcal{B}$-phase; *(iii)* if Sender did adopt a truth telling strategy, the expected cost of falling into a $\mathcal{B}$-phase is almost 0.

An important feature of the trust mechanism is that it keeps Receivers uncertain about when the Sender is next vulnerable to a review change. To see why this aids the Sender's ability to earn rents, consider the Firm's optimal strategy (from Section 2) in a Review Mechanism. At time 0, the Firm knows that it will have the veracity of its reports tested statistically T periods later. However, when T is large enough for the Law of Large Numbers to make these tests reliable, the marginal effect of a single misreport on the chances of failing are insignificant.[23] Thus, the Firm has an incentive to make some misreports, and due to discounting, it will prefer to overreport 'Good' at $t = 0$. But, if the first Mayor knows she is first in line, she will be aware of the Firm's incentive and treat it as babbling, undermining its incentive to 'overreport' in the first place. This incentive to front-load misreports causes the Firm's ability to extract persuasive rents to unravel when Mayors know their position in the line.

However, when aggregation systems leave Mayors uncertain about when the Firm is up for its next review, this uncertainty considerably restricts the Mayor's skepticism on seeing an 'Accept' recommendation. Then, despite the incentive of the Firm to front-load its lies, the Mayor can do no better than to follow the recommendations since the Review Mechanism is designed to keep the Sender 'close to' truthful on average across time.

6 Emergent Communication Mechanisms

The role of the CnC mechanism is to allow Receivers to allow the Sender to mix from the perspective of the Receiver but allows the Receiver to subsequently verify that the Sender was mixing according to some particular rule. This allows the Receiver to only punish the Sender off path. There are two classes of game where these properties endogenously emerge in equilibrium: 1) one state may simultaneous actions; and 2) many states many simultaneous actions. We will illustrate this first case by adapting the repeated public goods game developed in Hermalin [2007], this simultaneously illustrates a setting where sustaining greater levels of persuasion can be socially optimal. In the second example we use a finance case which relates to some real world attempts to improve persuasion and also speaks to how we can do better when there is noisy feedback about what the sender observes.

[23]The effect of a single misreport on the average is of the order of T^{-1}. By contrast, to ensure truth-telling passes the test with a high probability, the test's 'error bounds' are of the order $T^{-\frac{1}{2}}$.

6.1 One State Many Agents: A Repeated Public Goods Game

Consider an even numbered team of $L>3$ myopic workers (receivers) and one patient benevolent boss (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 $\theta_t \in \{1,3\}$ where $Pr(\theta_t = 3) = \mu_0 = \frac{1}{3}$ 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

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

The bosses stage payoff will just the average payoff of the workers, i.e., the boss, like the workers, wants everyone to exert effort independent of the state. 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 $\{\Pi_0, \Pi_1, ..., \Pi_{t-1}\}$.

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 the workers are myopic they cannot enforce high effort through punishment strategies.

If the boss is only able to send a single 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 individualized 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. This equilibrium Pareto dominates truth telling.[24]

Note that achieving higher levels of persuasion does not rely on this being a public good. For example, if their were no public good problem and the workers above just earned $\theta - e^l$ while the boss gets a payoff of $\sum e^l$ the same equilibrium as described above is still feasible. 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. So long as the sender's choice of *to whom* to send a message is ex-ante stochastic at the stage game from the receivers perspective we can get obedient behaviour from the receivers while still being able to verify that the sender is sticking to the equilibrium strategy ex-post - just as with the CnC mechanism.[25]

A consideration in these cases is the incentives and capabilities of Receivers to share their messages with one another at the stage game. In the public good game examined above the workers would never want to discourage others from working, and so it is plausible that even if able they would not have meaningful communication with one another. However, if we take the variant where there is no public good the workers have weak incentives to communicate with one another at the stage game; in this case, and similar, the sender would find it desirable to prevent communication between receivers, such as Chinese walls or departmental segregation.

6.2 Many States and Many Agents: Financial Advice & Disclosure Rules

We will illustrate this using an example of brokerages (sender) giving investment advice to their clients (receivers). There are several papers showing that advisers and 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]. This could be modeled using an asset pricing model with costly transactions and information acquisition, but this would be a paper in itself. So for expositional purposes we use a simple binary action model where a brokerage advises clients on products that it always wants them to buy.

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

[25]If the sender were not indifferent this would be okay so long as the sender's did not know who the sender had

There are L investors, each investor l is making a decision about an idiosyncratic portfolio that, given their preferences, is overvalued or undervalued $\theta_t^l \in \{Low, High\}$. The $Pr(\theta_t^l = High) = \mu_0 = 1/3$ for all t. [26] The portfolio values are independent across investors, $E[\theta_t^l|\theta_t^{k\neq l}] = E[\theta_t^l]$. [27] Let investor payoffs be identical to the mayor and firm case so that the investor buys if $\mu_t^l \geq 0.5$; The brokerage always wants their clients to buy. The brokerage observes θ_t^l for all l and sends a vector of messages $\boldsymbol{\mu}_t = (\mu_t^1, ..., \mu_t^L)$ to investors who then choose to buy or sell.

If investor l only observes the history of their own portfolio l and the respective advice given to them then we just have L replications of the standard case and the equilibrium payoff set is identical to that of the mayor and firm example. However, if investors can observe some measure of the history of all messages then we can get arbitrarily close to the set of payments available under full commitment. For concreteness, suppose investors observe the proportion of 'buy' recommendations sent by the brokerage in the past in addition to their own advice. Now the proportion of buy recommendations can be used in a fashion that is similar to the CnC. The brokerage is given a per period budget of 'buy' offers to allocate across investors and tells investors to buy when $\theta_t^l = High$ and then gives the remaining buy recommendations at random to the remaining investors. If the budget is less than or equal to two thirds then it is incentive compatible for the investors to follow the advice. This is sustained, of course, by the threat of moving to a babbling equilibrium. [28]

Moving from a world in which investors only observe their own history of messages to observing the aggregate of messages makes a far more profitable equilibrium for the brokerage. It so happens, that in 2002, the National Association of Security Dealers, 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] analyse 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. On introduction of these new rules Buy recommendations dropped almost immediately to 51% and by the following year Buy recommendations were down to 42%. This was not a drop from an anomalous high, since the beginning of the data they analyse (1996) Buy

[26] We could relax this assumption. It is sufficient just to have ex-post realisation of the overall success rates across assets.

[27] This assumption is substituting for a model where the investor gives advice on L asset classes, there are many investors each with an idiosyncratic component to their portfolio, and investors are not contemporaneously monitoring the advice given on all assets. Such a model can easily be grounded in rational inattention or costly information acquisition. Further, we know as an empirical matter, that different investors do in fact focus on different asset classes.

[28] The lack of competition may be of concern, we know however from Gentzkow and Kamenica [2016] that there is still a role for persuasion in competitive environments. But, we are assuming that something akin to this result persists in a dynamic environment without commitment.

recommendations had always exceeded 60%. Along with the drop in buy recommendations Kadan et al. [2009] show that prices became more responsive to buy recommendations, suggesting investors found the new reports more persuasive.

This is compatible with our model. Suppose that prior to the change the brokerages were recommending buy when the assets were undervalued and some proportion of the time when the assets were overvalued. On path they would be punished going through babbling periods with investor l where the firm always says accept and the investor ignores the brokerage. [29]

Using an adaptation of our model, we can study how the introduction of these disclosure rules affected the ability of brokerage firms to persuade clients to trade with them. In the absence of these rules, analysts faced repeated interaction with clients. Given the results of Section 4, we would expect any informative communication in this setting to be sustained by long periods of babbling as a clients punish overly optimistic recommendations. Indeed, the aggregate reporting data bear this out somewhat - reports were skewed heavily towards 'Buy' calls, and yet simultaneously had little effect on market behaviour. Nonetheless, the nature of repeated interaction may have provided enough discipline to prevent analysts from completely babbling.

One effect of the rule change was to allow investors to see the distribution of recommendations made by the investment bank's analysts across assets, providing context for the recommendation made in an individual report. In doing this, the NASD rules effectively provided banks with a commitment device to enact optimal persuasion. Indeed, we show that the introduction of aggregate reporting standards introduces equilibria in which the banks can achieve the optimal payoffs from persuasion. In equilibrium, the aggregate reporting standard can be used as a payoff-irrelevant disciplining device, around which investors can 'punish' the bank with babbling if its aggregate reports become excessively skewed toward 'Buy' (when compared to individual outcomes, ex post). Given such a device and a large enough asset space, an equilibrium can be sustained in which the bank credibly reports according to the optimal persuasive strategy and in which Receivers never need to employ on-path punishments, since the aggregate statistics can be used as a disciplining device. Notice that, from the perspective of a single asset, we have introduced a payoff-irrelevant 'state' variable which is seen first by the bank and only later by investors. - this is simply the vector of performance outcomes across the other assets.

[29] This form of babbling is even more natural in a world with naive investors who take advice at face value.

7 Conclusion

There are three main results in this paper. The first, in line with previous thinking, is that meaningful communication today can be sustained out of the desire for meaningful communication in the future. The second, is that while repeated interaction can generate commitment to dishonest reporting strategies it cannot generally achieve the optimal outcome for the Sender. The third is that there are various alterations to the game that achieve the optimal outcome for the Sender.

These results have important implications for the real world, and the literature on persuasion. Perhaps the most important implication is that the use of a simple mechanism such as a CnC can generate Pareto improvements for Senders and Receivers. Further, mediators can drastically improve outcomes by serving to aggregate reports and reduce the necessity of on the path punishments in the presence of noise. However, it also follows that the introduction of such mechanisms in some circumstances can increase the extent to which dishonest persuasion is possible, and give Senders an incentive to push the game towards less honest equilibria.

It seems that some of these communication mechanisms have already been set up to some extent either through the use of laws, third parties or online technologies. Although, it is not clear the extent to which they've been designed to be optimal from the point of view of the Senders, Receivers, or society. Finally, it seems that not all tools are currently in use, and that it might be possible to benefit both the Senders and Receivers of the world by identifying those cases where they can be fruitfully implemented.

References

S. Athey and K. Bagwell. Optimal collusion with private information. *RAND Journal of Economics*, 32(3), 2001.

S. Athey and K. Bagwell. Collusion with persistent cost shocks. *Econometrica*, pages 493–540, 2008.

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.

D. Barron. Attaining efficiency with imperfect public monitoring and one-sided markov adverse selection. *Theoretical Economics*, 2016.

J. Bizzotto, J. Rüdiger, and A. Vigier. Delayed persuasion. Technical report, University of Oxford Working Paper, 2016.

A. Chakraborty and R. Harbaugh. Persuasion by cheap talk. *American Economic Review*, 100(5):2361–2382, 2010.

V. Crawford and J. Sobel. Strategic information transmission. *Econometrica*, 50(6):1431–1451, 1982.

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. Working Paper., 2015.

J. F. Escobar and J. Toikka. Efficiency in games with markovian private information. *Econometrica*, 81(5):1887–1934, 2013.

M. Gentzkow and E. Kamenica. Competition in persuasion. *The Review of Economic Studies*, (September):1–36, 2016. URL `http://faculty.chicagobooth.edu/matthew.gentzkow/research/BayesianComp.pdf{%}5Cnhttp://www.nber.org/papers/w17436`.

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.

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.

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.

C. Margaria and A. Smolin. Dynamic Communication with Biased Senders. 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 Long-Run Uncertainty. 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.

J. Renault, E. Solan, and N. Vieille. Dynamic sender–receiver games. *Journal of Economic Theory*, 148(2):502–534, 2013.

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, 2015.

Appendix A

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

Assumption. *For any $\mu_1, \ldots, \mu_n \in \Delta_P \Theta$ such that $v(\mu_1) = v(\mu_2) = \cdots = v(\mu_n)$, where $\mu_i \neq \mu_j$ for at least two $i, j \in \{1, 2, \ldots, n\}$, $n > 1$, the concavification of v on Δ_P, $\hat{v}_P(\mu)$, satisfies*

$$\hat{v}_P\left(\sum_{i=1}^{n} \alpha_i \mu_i\right) > \sum_{i=1}^{n} \alpha_i v(\mu_i)$$

Proof of Lemma 1

Proof. Take any equilibrium of game Γ^∞, and denote the corresponding strategies and beliefs respectively by $\sigma^\star = (\sigma_t^\star(\phi_t, \theta_t))_{t=0}^{\infty}$, $\rho_t^\star(\phi_t, m_t)$, $t = 0, 1, 2, \ldots$, for S, R_t, and $\mu_t^\star(\phi_t, m_t)$, $t = 0, 1, 2, \ldots$. Since these strategies and beliefs form an equilibrium of Γ^∞, they obey conditions 1. - 3. in section 3.2. 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. in section 3.3.

We assume for ease of exposition that there is a message $\underline{m} \in M$ which is never played on the equilibrium path.[30] For each public history $\phi_t = (m_\tau, a_\tau, \theta_\tau)_{\tau=0}^{t}$, define the function $\tilde{\mu}_t(\phi_t, m_t) := \mu_t^\star(\phi_t, m_t)$ and a corresponding history of game $\hat{\Gamma}^\infty$, $h_t^\phi := (\tilde{\mu}_t(\phi_t, m_t), a_t, \theta_t)$. Let the set of all such histories be H^ϕ. For any history h_t of game $\hat{\Gamma}^\infty$ such that $h_t \notin H^\phi$, recursively define ϕ_t' as a history of Γ^∞ in which $m_\tau = \underline{m}$ for any τ such that $\tilde{\mu}_t \neq \tilde{\mu}_t(\phi_\tau', m_\tau)$ for some $m_t \in M$. Consider now strategies $\hat{\sigma}$, $\hat{\rho}_t$, $t = 0, 1, \ldots$, and beliefs $\mu_t(h_t, \tilde{\mu}_t)$ in game $\hat{\Gamma}^\infty$, where:

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

where $\lambda_t^\star \in \Delta(\Delta\Theta)$ satisfies $\Pr(\hat{\mu} = \tilde{\mu}_t(\phi_\tau', m_\tau) \mid \phi_t, \theta_t) = \sigma_t^\star(m_\tau; \phi_t, \theta_t)$ and $\sigma_t^\star(m_\tau; \phi_t, \theta_t) := \Pr(m = m_\tau \mid \phi_t, \theta_t; \sigma_t^\star(\phi_t, \theta_t))$ is the measure over M induced by lottery $\sigma_t^\star(\phi_t, \theta_t)$,

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

[30]Lemma 3.3 does not require such an assumption (proof available on request), but it significantly reduces the notation required to describe strategies fully.

and

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

where $a_t=\underline{a}\in A$ is the action which minimizes S's stage payoff.

At any history $h_t^{\phi}\in H^{\infty}$, strategy $\hat{\sigma}$ assigns probability $\sigma_t^{\star}(m_\tau;\phi_t,\theta_t)$ to message $\hat{\mu}_t=\mu_t(\phi_t,m_t)$. Moreover, given message $\hat{\mu}_t$, $\hat{\rho}_t\left(h_t^{\phi},m_t\right)$ induces the same lottery over $a\in A$ that $\rho^{\star}(\phi_t,m_t)$ 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 $(\sigma^{\star},\rho^{\star})$ does in Γ^{∞}. Moreover, the transition probabilities between h_t^{ϕ} and $h_{t+1}^{\phi}=\left(h_t^{\phi},\tilde{\mu}_t(\phi_t,m_t),a_t,\theta_t\right)$ are clearly identical to those between ϕ_t^{ϕ} and $\phi_{t+1}^{\phi}=\left(\phi_t^{\phi},,m_t,a_t,\theta_t\right)$. Thus, $(\hat{\sigma},\hat{\rho})$ induces the same ex ante distribution over a_t as $(\sigma^{\star},\rho^{\star})$ as well.

Given S's strategy, a message $\hat{\mu}_t=\tilde{\mu}_t(\phi'_\tau,m_\tau)$ induces R_t's beliefs at any history $h_t^{\phi}\in H^{\phi}$ to be

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

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.

Finally, we argue that 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 $\hat{\mu}_t=\tilde{\mu}_t(\phi'_\tau,m_\tau)$ sent by S, since $\rho^{\star}(\phi_t,m_t)$ 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 $\hat{\mu}_t=\tilde{\mu}_t(\phi'_\tau,m_\tau)$, for some τ, $m_\tau\in M$, since both such strategies are relabelings of $\sigma^{\star}$ and some alternative feasible strategy σ' in game Γ, respectively (and similarly for $\hat{\rho}$, $\rho^{\star}$). Since $\sigma^{\star}$ is optimal in game Γ^{∞}, $\hat{\sigma}$ is preferred to $\hat{\sigma}'$. Finally for any other deviation, we have specified the strategies for Sender and Receiver in such a way that the stage payoff from 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 τ, $m_\tau\in M$, and moreover, the continuation payoff would be the same as under $\hat{\sigma}'$. $\square$

Proof of Theorem 1

Proof. Fix μ_0. Let U be the set of all real-valued upper-semicontinuous functions on $\Delta\Theta$, 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 $U^\star$, defined as

$$U^\star(\mu_0) = \left\{ v \in U: \ \hat{v}(\mu_0) > \sum_i \alpha_i v(\mu_i), \forall \{\mu_i\}_{i=1}^{N+1}, s.t. \, \mu_0 = \sum_i \alpha_i \mu_i, v(\mu_i) = v(\mu_j), \mu_i \neq \mu_j, i, j = 1, 2, \right.$$

is open and dense.[31]

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

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

where $\tilde{\Lambda}(\mu_0)$ is the subset of $N+1$-point distributions in $\Lambda(\mu_0)$.[32] Since v' is upper semi-continuous and $\tilde{\Lambda}(\mu_0)$ is compact in λ, the function $\sum_{i=1}^{N+1} \lambda_i v(\mu_i)$ attains its maximum in $\tilde{\Lambda}(\mu_0)$.[33] Moreover the set $\tilde{\Lambda}(\mu_0)$ is both upper and lower hemicontinuous in μ_0.To establish upper hemicontinuity, take a sequence $\{\mu_0^n\} \to \mu_0$, and any corresponding convergent sequence, $\{\lambda^n\} \to \lambda$, $\lambda^n \in \Lambda(\mu_0^n)$. We show that $\lambda \in \Lambda(\mu_0)$. Suppose not. Then, for any $\epsilon > 0$, there always exist a subsequence $\{n'\} \in \mathbb{N}$ such that

$$\left| \sum_{i=1}^{N+1} \lambda_i^{n'} \mu_i^{n'} - \mu_0 \right| > \epsilon$$

$\forall n'$. But since $\{\mu_0^n\} \to \mu_0$, and $\left|\lambda_i^{n'}\right| < 1$, $\forall i$, n', for some $0 < \epsilon' < \frac{\epsilon}{N+1}$, there must exist an m such that for all $n' \geq m$, we have

$$\left| \sum_{i=1}^{N+1} \lambda_i^{n'} \mu_i^{n'} - \mu_0 \right| \leq \sum_{i=1}^{N+1} \lambda_i^{n'} \left| \mu_i^{n'} - \mu_0 \right| < \epsilon$$

- a contradiction. To show lower hemicontinuity, take any sequence $\{\mu_0^n\} \to \mu_0$, and any

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

[32] By Caratheodory's Theorem, this is without loss for finding $\hat{v}'$.

[33] Indeed, it is easy to see that $\tilde{\Lambda}(\mu_0)$ is a compact subset of $\mathbb{R}^{N(N+2)}$.

$\lambda \in \Lambda(\mu_0)$. We show that there exists a subsequence $\{\mu_0^{n''}\}$, such that $\exists \lambda^{n''} \in \Lambda(\mu_0^{n''})$ satisfying $\lambda^{n''} \to \lambda$. Fixing some $\epsilon > 0$, we can find vertices $p(\epsilon) = (\mu_1(\epsilon), \ldots, \mu_N(\epsilon))$ for which all $\tilde{\mu}_0$ satisfying $|\tilde{\mu}_0 - \mu_0| \leq \epsilon$ can be expressed as convex combinations of the vertices in$p(\epsilon)$ *(Rockafeller, Theorem 20.4)*. Moreover for an $\frac{\epsilon}{2^k}$-ball, $k \in \mathbb{N}$, $k \geq 1$, we can enclose all points in $|\tilde{\mu}_0 - \mu_0| \leq \frac{\epsilon}{2^k}$ by the simplex generated by $p_k(\epsilon) = \left(\frac{2^k-1}{2^k}\mu_0 + \frac{1}{2^k}\mu_i(\epsilon)\right)_{i=1}^{N}$. Noting that we can write $\mu_0 = \sum \lambda_i \mu_0^i$, each vertex of $p_k(\epsilon)$ can be written $\sum \lambda_i \mu_0^i + \frac{1}{2^k}(\mu_i(\epsilon) - \mu_0)$. Thus, for any sequence $\mu_0^{n''} \to \mu_0$ and $m \in \mathbb{N}$, we can find k such that all $\mu^{n''}$ can be expressed as a convex combination of points p_k, which get arbitrarily close to $\sum \lambda_i \mu_0^i$. This establishes the limiting sequence $\lambda^{n''} \to \lambda$.

Applying Berge's Theorem of the Maximum, the value function, $\hat{v}'$, is continuous on $\Delta\Theta$. Consequently, the subgraph of $\hat{v}'$, $sub(\hat{v}') = \{(\mu, \nu) : \nu \leq \hat{v}'(\mu), \mu \in \Delta\Theta\}$, is a closed convex set. Bound $sub(\hat{v}')$ 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, $H(\hat{v}') := sub(\hat{v}') \bigcap \{(\mu, v) : v \geq \underline{B}\}$. Note that $H(\hat{v}' + \epsilon)$ has the same properties and contains $H(\hat{v}' + 2\epsilon)$.

We are now able to find an ϵ-perturbation of v' such that the new function v satisfies $v \in U^\star$. Partition $\Delta\Theta$ into two sets: $C = \{\mu : v'(\mu) = \hat{v}'(\mu)\}$ and $\overline{C} := \Delta\Theta / C$. We now construct a polyhedral convex set P for which $int(sub(\hat{v})) \subset P \subset int(sub(\hat{v} + \epsilon))$ and all the vertices of P, $\{(\mu_i^\star, \nu_i^\star)\}_{i=1}^{M}$, for which $\nu_i > \hat{v}'(\mu)$ satisfy $proj_{\Delta\Theta}(v_i) \in C$. For any $x \in H(\hat{v}')$, we can choose a simplex S_x such that $x \in S_x$ and $S_x \in int(D)$. Because $\sum_{i=1}^{N+1} \lambda_i v(\mu_i)$ attains its maximum on $\Delta\Theta$, we can in fact choose S_x such that its vertices $(s_1, s_2, \ldots, s_{N+1})$ satisfy $s_i = (\mu_i', \hat{v}'(\mu_i) + \epsilon)$ and $\mu_i' \in C$, if $s_i \notin int(H(\hat{v}))$. From this union of simplices $\bigcup \{S_x\}$, we can find a finite subset of simplices whose convex hull also covers $H(\hat{v}')$ - this is the polyhedron P (Rockafeller, Theorem 20.4). Moreover, by construction, no vertex of P that lies above $H(\hat{v}')$ has $\mu_i^\star \in \overline{C}$.

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

We now show that $U^\star$ is open in v, for all priors close to μ_0, for any μ_0 satisfying $v(\mu_0) < \hat{v}(\mu_0)$. Specifically, we show that for any function $v \in U^\star$, there exist ϵ_1, $\epsilon_2 > 0$ s.t. for all $\tilde{v}$ satisfying $||\tilde{v} - v|| < \epsilon_1$, $\tilde{\mu}_0$ satisfying $|\tilde{\mu}_0 - \mu_0| < \epsilon_2$, we have $\tilde{v} \in \bigcup_{|\tilde{\mu}_0 - \mu_0| < \epsilon_2} U^\star(\mu_0)$.

Take some $v \in U^\star$. We argue that, for 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 $|\hat{v}(\tilde{\mu}_0) - \sum_i \alpha_i v(\mu_i)| < \delta_1$, for some $\lambda \in \Lambda(\tilde{\mu}_0)$ then $|v(\mu_i) - v(\mu_j)| \geq \tilde{\epsilon}_2$

for some $i \neq j$, $i,j \in \{1,2,\ldots,N+1\}$.[34] Suppose this were not the case. Then, for any δ, $\tilde{\epsilon}_1$, $\tilde{\epsilon}_2 > 0$, 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)* $|v(\mu_i) - v(\mu_j)| < \tilde{\epsilon}_2$, *(iii)* $|\hat{v}(\tilde{\mu}_0) - \sum_i \alpha_i v(\mu_i)| < \delta_1$. Now consider any sequence $(\delta^n, \tilde{\epsilon}_1^n, \tilde{\epsilon}_2^n)_{n=1}^{\infty}$ satisfying $\lim_{n\to\infty}(\delta^n, \tilde{\epsilon}_1^n, \tilde{\epsilon}_2^n) = 0$. 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 $\left((\tilde{\mu}_0^{n'}, \lambda_0^{n'})\right) \to (\mu_0, \lambda)$ for some $\lambda^\star \in \Lambda(\mu_0)$.[35] Moreover, upper semi-continuity of v implies that at this limit, we must either have *(i)* $\hat{v}(\mu_0) = \sum_i \lambda_i^\star v(\mu_i^\star)$, and $v(\mu_i) = \hat{v}(\mu_j)$, $\forall i,j \in \{1,2,\ldots,N+1\}$; *(ii)* $\sum_i \alpha_i v(\mu_i) > \hat{v}(\mu_0)$, or *(iii)* $\hat{v}(\mu_0) = \sum_i \lambda_i v(\mu_i)$, $v(\mu_i) \neq v(\mu_j)$, for at least two i,$j \in \{1,2,\ldots,N+1\}$.[36] Since $v \in U^\star$, case *(i)* yields a contradiction. By definition of $\hat{v}$, case *(ii)* also implies a contradiction. Finally, we rule out case *(iii)*. Since there the discrete upward jump at $\mu_i^\star$, for some i, must also cause a discontinuity at $\hat{v}(\mu_0)$ on the path $\mu_0^n \to \mu_0$ - a contradiction, to the continuity of $\hat{v}$, which we proved above.

Finally, for any perturbed function v' such that $||v' - v|| \leq \min\left\{\frac{\delta}{3}, \frac{\epsilon_2}{3}\right\}$, we must also have $\forall\, |\tilde{\mu}_0 - \mu_0| < \tilde{\epsilon}_1$, if $|\hat{v}'(\tilde{\mu}_0) - \sum_i \alpha_i v'(\mu_i)| < \delta_1'$, for some $\lambda \in \Lambda(\tilde{\mu}_0)$ then $|v'(\mu_i) - v'(\mu_j)| \geq \tilde{\epsilon}_2'$ for some $i \neq j$, $i,j \in \{1,2,\ldots,N+1\}$, for some $\tilde{\epsilon}_2'$, $\delta_1' > 0$. □

Proof of Lemma 2

Proof. 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 $M' = \cup_{\theta \in \Theta_t} supp(\sigma_\tau(h_\tau, \theta))$. This strategy induces a N'-point distribution $\nu \in \Delta\left(\Delta\Theta^{N'}\right)$ 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

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

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

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

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

[35] A variation on our argument that $\Lambda(\mu_0)$ is compact in λ can be used to establish compactness in (μ_0, λ).

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

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

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

Given any history, we define an equilibrium message $m^\theta \in M'$ to be *uniquely proscribed* at state θ if $supp\left(\sigma_\tau\left(h_\tau, \theta\right)\right) = \left\{m^\theta\right\}$. The set of all messages that are uniquely proscribed 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 proscribed, $m \in M^\Theta$, and those that are not, $m \in M'/M^\Theta$.

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

$$\sum_{m_\tau \in M'/\{\tilde{m}\}} \alpha_{m_\tau} \mu_\tau\left(m_\tau\right) = \mu_0 \tag{6}$$

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, $\Delta\left(\Delta\Theta^{N'}\right)$). By Proposition 1 in Kamenica & Gentzkow (2011), the posteriors $\mu_\tau\left(\tilde{m}\right) \in \left\{\mu_\tau\left(m\right)\right\}_{m \in M'/\{\tilde{m}\}}$ can be sustained by a feasible signal structure with $N'-1$ distinct messages. Moreover, the message $\tilde{m}$ cannot be uniquely proscribed in any state $\theta \in \Theta$. Otherwise, there would exist some θ^i for which $\mu_\tau^i\left(m\right) = 0$, $\forall m \in M'/\left\{\tilde{m}\right\}$, while $\mu_0^i > 0$, violating (6). Therefore, $\tilde{m} \in M'/M^\Theta$ and for every state θ in which σ proscribes $\Pr\left(m_\tau = \tilde{m} \mid h_\tau, \theta\right) > 0$, there exists another message m'_θ sent with positive probability in state θ.

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

Therefore, strategy $\sigma^\star$ achieves the same payoffs for the Sender from history h_τ, 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 $\mathbb{E}\left[V\left(h_\tau, \theta_\tau\right)\right]$. $\square$

Proof of Proposition 1

Proof. Let v^B be the worst expected stage payoff to the sender from any equilibrium of the stage game. Define σ_t^B as the equilibrium strategy of the stage game that give the Sender v^B, given the Receiver's equilibrium beliefs. At any stage t, $\sigma_{t+\tau}^B$ for all τ constitutes a continuation equilibrium. Define σ_t^T as the stage game strategy of reporting the true state of the world $\tilde{\mu}_t^i = 1$ if and only if $\theta_t = \theta_t^i$ and 0 otherwise. Let v^T be the expected stage payoff to the Sender from a stage game conditional on the sender following the strategy σ_t^T. Let $H_t^T \subset H_t$ be the set of histories at time t consistent with the sender following strategy σ_τ^T at all $\tau < t$. Define the following strategy σ^T such that for all t:

$$\sigma_t = \begin{cases} \sigma_t^T & \text{if } h_t \in H_t^T \\ \sigma_t^B & \text{if } h_t \notin H_t^T \end{cases}.$$

Given such a strategy and the obedient beliefs of the Receiver, the gap in the continuation payoffs from playing the strategy σ_t^T at a history $h_t \in H_t^T$ is:

$$\frac{\delta(v^T - v^B)}{1-\delta}.$$

At any history $h_t \in H_t^T$ the stage payoff from deviating is bounded above by $\bar{v} < \infty$, as the Senders stage payoff is upper semi-continuous and the message space is compact. Hence, if $v^T > v^B$, there exists $\bar{\delta} < 1$ such that

$$\frac{\delta(v^T - v^B)}{1-\delta} > \bar{v} \text{ for all } \delta > \bar{\delta}.$$

σ^T constitutes an equilibrium strategy with obedient Bayes plausible beliefs for $\delta > \bar{\delta}$. □

Proof of Proposition 2

Proof. In any equilibrium, S must be indifferent at any history (h_t, θ_t^i) between all messages $\tilde{\mu} \in \sigma_t(h_t, \theta_t^i)$. Since $\underline{\mu}_i(h_t) := \arg\min \underline{v}_i(\sigma_t(h_t, \theta_t^i))$ 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

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

Consider the following problem:

$$\sup_{\sigma \in \Sigma} \mathbb{E}_\theta \left[V_t\left(\underline{h}_t, \theta_0\right)\right] \tag{7}$$

s.t.

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

$$\mu_0 \in co\left(\cup_{\theta_t\in\Theta} supp\left(\sigma_t\left(h_t,\theta_t\right)\right)\right),$$

$\forall \underline{h}_t$ s.t. $\tilde{\mu}_t = \underline{\mu}^i_{\tau}$ at all subsequences $\underline{h}_{\tau'}$, $0 \leq \tau' \leq \tau$, of $\underline{h}_t$ at which S acts. We refer to the set of continuation payoffs that satisfy all constraints in (7), by $\mathcal{V}$. Notice $\mathcal{V}$ is non-empty.[37]

At $t = 0$ (where $h_0 = \emptyset$), problem (7) is a relaxed version of problem (4): it only retains constraints for histories in which S has always reported the 'worst' current message $\underline{\mu}^i_{\tau'}$ among all those available in the support of his strategy at previous histories, $h_{\tau'}$, $\tau' < \tau$. All other constraints from (4) are dropped. Thus, the optimal value of (7) provides an upper bound on (4).

Let $V^\star_t\left(\underline{h}_t,\theta_t\right)$ be the supremum achieved in problem (7) at history $\underline{h}_t$. From the first constraint, we must have

$$\mathbb{E}\left[V^\star_{t+\tau}\left(\underline{h}_{t+\tau},\theta_{t+\tau}\right)\right] = \sup_{\sigma_t(h_t,\theta),V_{t+\tau+1}} \mathbb{E}\left[v\left(\underline{\mu}^i_{t+\tau}\right)+\delta V_{t+\tau+1}\left(\left(\underline{h}_{t+\tau},\underline{\mu}^i_{t+\tau},\theta_{t+\tau}\right),\theta_{t+\tau+1}\right)\right] \tag{8}$$

where the supremum is taken over feasible lotteries $\sigma_t\left(h_t,\theta\right) \in \Lambda\left(\mu_0\right)$ and feasible payoffs from the continuation equilibrium, $V_{t+\tau+1} \in \mathcal{V}$.[38]

Notice that, for any $t+\tau$, history $\left(\underline{h}_{t+\tau},\theta_{t+\tau}\right)$ and corresponding $\sigma_{t+\tau}\left(\underline{h}_{t+\tau},\theta_{t+\tau}\right)$, $V_{t+\tau}\left(\underline{h}_{t+\tau},\theta_{t+\tau}\right)$ is maximized by choosing the highest feasible expected continuation, $\mathbb{E}\left[V^\star_{t+\tau+1}\left(\left(\underline{h}_{t+\tau},\underline{\mu}_{t+\tau},\theta_{t+\tau}\right),\theta_{t+\tau+1}\right)\right]$ Moreover, since the continuation games at histories $\left(\underline{h}_{t+\tau}\right)$ and $\left(\left(\underline{h}_{t+\tau},\underline{\mu}_{t+\tau},\theta_{t+\tau}\right)\right)$ are identical, the expected continuation values must be equal:

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

Substituting into (8) yields, on rearrangement:

$$\left(1-\delta\right)\mathbb{E}_\theta\left[V^\star_0\left(\theta_0\right)\right] = \max_{\lambda\in\Lambda(\mu_0)} \sum \mu^i_0 \underline{v}_i\left(\lambda\right) \tag{9}$$

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

Since $\underline{v}_i\left(\lambda\right)$ is the minimum of finitely many upper semicontinuous functions (from Lemma 2, we need only choose from N distinct posteriors, μ, and v is upper semicontinuous), it is

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

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

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 $\lambda^\star \in \Lambda(\mu_0)$, with associated support $\{\mu_1^\star, \mu_2^\star, \ldots, \mu_{N'}^\star\}$, where for convenience we index such that $v(\mu_1^\star) \leq v(\mu_2^\star) \leq \cdots \leq v(\mu_{N'}^\star)$.

Following Proposition 1, v^B is the worst expected stage payoff to the sender from any equilibrium of the stage game; σ_t^B is the equilibrium strategy that induces v^B at any stage t; $\sigma_{t+\tau}^B$ is a feasible continuation equilibrium for all $1 < \tau' \leq \tau$ and any τ. Now define $\underline{v}^* := \sum_{i=1}^N \mu_0^i \underline{v}(\lambda^*)$, this is the putative upper bound on the average stage payoff described in Proposition 2.Now define σ_t^* as the stage game strategy of inducing the lottery λ^*. Consider the following strategy σ^* such that for all t:

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

Where H_t^B is the set of histories inconsistent with σ^* at t. Hence, if $h_t \in H_t^B$ then $h_{t+\tau} \in H_{t+\tau}^B$ for all $\tau > 0$. Given strategy σ^* we only play at histories in $\underline{H}_t$ or $\bar{H}_t$. At the beginning of the game $h_0 \in \underline{H}_0$. If $h_t \in \bar{H}_t$ we say that we are in a punishment period. A punishment period commences if $\tilde{\mu}_t \neq \underline{\mu}^i$ at $h_t \in \underline{H}_t$ for some $\tilde{\mu}_t \in supp(\lambda^\star)$ and lasts for $K_{i,\mu} + \mathbf{1}(b_{t+1} < \beta_{i,\mu})$ periods. $K_{i,\mu}$ and $\beta_{i,\mu} \in [0,1]$ are the values of $k$and β that solve the following equation:

$$v(\mu) - v(\underline{\mu}_i) = \beta\delta^{k+1}(\underline{v}^* - v^B) + \sum_{\tau=1}^{k} \delta^\tau(\underline{v}^* - v^B). \tag{10}$$

First, note that if $\underline{v}^* = v^B$ then this payoff can be sustained as an average payoff in an equilibrium of the stage game and hence also in the repeated game, leaving us to focus on the case where $\underline{v}^* > v^B$. When $\underline{v}^* - v^B > 0$ then let $\underline{\delta} < 1$ be defined by the following equation:

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

Hence, for $\delta > \underline{\delta}$ there always exists a $k$and β solving equation 10 as $v(\mu) - v(\underline{\mu}_i) < \bar{v} < \infty$. Now, to see that σ^* is an equilibrium for $\delta > \bar{\delta}$, first note that σ_t^B is an equilibrium of the stage game and so can be an equilibrium at any history; second, there is a solution to equation (10) and so the Sender is made indifferent between choosing $\tilde{\mu}_t$ and $\underline{\mu}^i$ for all i and all $\tilde{\mu}_t \in supp(\lambda^\star)$; and finally, deviating from σ_t^* for some $h_t \in \underline{H}_t$ results in an expected utility loss of

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

$\square$

Proof of Theorem 2

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 (2), 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 $M := \{m_1, m_2, \ldots, m_{N'}\}$, where $N' \leq N$. 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 $m^{-1}(m_j) \cap m^{-1}(m_k)$, $\forall j, k \in \{1, 2, \ldots, N'\}$, $j \neq k$ and $\cup_{j \in \{1,\ldots,N'\}} m^{-1}(m_j) = \Theta$. Under 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 j^{th} entry of μ is

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

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

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

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

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

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 2 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 the discounted average payoff from optimal signal design on each

Receiver, $\hat{v}(\mu_0)$, cannot be obtained by any partition strategy. Take any optimal experiment $(M, s^{\star\star})$ that does achieve $\hat{v}(\mu_0)$, and denote the lottery over posteriors induced by the experiment by $\lambda^{\star\star} \in \Delta(\Delta\Theta)$. Let the support of this distribution be $\{\mu_1^{\star\star}, \mu_2^{\star\star}, \ldots, \mu_{N'}^{\star\star}\}$, and let the probability of posterior $\mu_j^{\star\star}$ under $\lambda^{\star\star}$ be $\lambda_j^{\star\star}$. Then, the expected payoff under lottery $\lambda^{\star\star}$ is

$$\sum_{j \in \{1,2,\ldots,N'\}} \lambda_j^{\star\star} v\left(\mu_j^{\star\star}\right) = \sum_i \mu_0^i \left(\sum_j \frac{\lambda_j^{\star\star} \mu_j^{i,\star\star}}{\mu_0^i} v\left(\mu_j^{\star\star}\right) \right)$$

where $\mu_j^{i,\star\star} := \Pr\left(\theta^i \mid \mu = \mu_j^{\star\star}\right)$ is the i^{th} component of vector $\mu_j^{\star\star}$, and $\sum_j \frac{\lambda_j^{\star\star} \mu_j^{i,\star\star}}{\mu_0^i} = 1$. However, by definition of $\underline{v}_i(\lambda^{\star\star})$ we have

$$\underline{v}_i(\lambda^{\star\star}) \leq \sum_j \frac{\lambda_j^{\star\star} \mu_j^{i,\star\star}}{\mu_0^i} v\left(\mu_j^{\star\star}\right) \tag{11}$$

We now argue that, generically, there must exist $i \in \{1, 2, \ldots, N\}$ such that (9) holds with strict inequality. Suppose not. Then experiment $s^{\star\star}$ must involve a partition of Θ into a set of subsets $\{P_1, P_2, \ldots, P_{N''}\}$ and corresponding partitions of M into $\{M_1, \ldots, M_{N''}\}$ such that

$$\Pr(m \in M_j \mid \theta \in P_k) \begin{cases} > 0 & , \ if\ j = k \\ = 0 & , \ otherwise. \end{cases}$$

$j, k \in \{1, 2, \ldots, N''\}$ and

$$v(\mu(m)) = v(\mu(m'))$$

$\forall m, m' \in M_j$, $j = 1, 2, \ldots, N''$. Since by assumption, this signal is not honest, there must exist at least one $j \in \{1, 2, \ldots, N''\}$ and messages $m, m' \in M_j$ such that $v(\mu(m)) = v(\mu(m'))$ but $\mu(m) \neq \mu(m')$. However, notice that these messages are only ever sent in the subset of $P_j \subset \Theta$. Therefore, for s** to maximize $\mathbb{E}[v(\mu)]$ on $\Lambda(\mu_0)$, it must also be maximizing $\mathbb{E}[v(\mu)]$ among all lotteries on $\gamma \in \Delta P_j$, subject to the restriction that

$$\sum_l \gamma_l \hat{\hat{\mu}}_l = \mu(P_j) \left(:= \Pr(\theta \mid \theta \in P_j)\right)$$

for some beliefs $\hat{\hat{\mu}}_l$ in the support of ΔP_j.[39] Refer to this feasible set as $\Gamma_{P_j}(\mu(P_j))$. However, by Theorem 1, this generically cannot be true: there exist arbitrarily small perturbations of v on the subset ΔP_j such that $v(\mu(m)) \neq v(\mu(m'))$ at the optimal value of $\mathbb{E}[v(\mu)]$ on the set $\Gamma_{P_j}(\mu(P_j))$.

[39]Note that $\hat{\hat{\mu}}_l^i = 0$ for $\theta^i \notin P_j$.

Thus, generically there must exist some state $\theta^i \in \Theta$ for which

$$\underline{v}_i\left(\lambda^{\star\star}\right) < \sum_j \frac{\lambda_j^{\star\star}\mu_j^{i,\star\star}}{\mu_0^i} v\left(\mu_j^{\star\star}\right)$$

and therefore

$$\sum \mu_0^i \underline{v}_i\left(\lambda^{\star\star}\right) < \sum_{j\in\{1,2,\ldots,N'\}} \lambda_j^{\star\star} v\left(\mu_j^{\star\star}\right) = \hat{v}\left(\mu_0\right) \tag{12}$$

A similar argument establishes that the payoff from *any* experiment inducing arbitrary lottery λ' attains a weakly higher payoff than $\sum \mu_0^i \underline{v}_i\left(\lambda'\right)$ evaluated at λ'. Therefore,

$$\sum \mu_0^i \underline{v}_i\left(\lambda^\star\right) < \hat{v}\left(\mu_0\right)$$

where $\lambda^\star$ solves (9). The inequality is generically strict since either: *(i)* $\lambda^\star$ is not an honest strategy (in which case (12) holds generically), or *(ii)* $\lambda^\star$ is honest, in which case $\sum_{j\in\{1,2,\ldots,N'\}} \lambda_j^\star v\left(\mu_j^\star\right) < \hat{v}\left(\mu_0\right)$ by our assumption that $\hat{v}\left(\mu_0\right)$ cannot be implemented by an honest experiment. □

Proof of Theorem 3

Proof. We first prove the result for $\nu^\star = \hat{v}\left(\mu_0\right)$. Defining the state variable as $\left(\theta_t, c_t\right) \in \Theta \times [0,1]$, it is easy to show that the optimal experiment under commitment can be written as honest with respect to $\left(\theta_t, c_t\right)$. The Theorem then follows as a Corollary of Theorem 2. Finally, for any $\underline{v}\left(\mu_0\right) \leq \nu^\star \leq \hat{v}\left(\mu_0\right)$, where $\underline{v}\left(\mu_0\right)$ is Sender's worst stage game payoff, we can achieve $\nu^\star$ 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 Grim Trigger to the worst stage equilibrium forever. □

Proof of Theorem 4

Proof. In arbitrary review phase $\mathcal{G}_j$, denote realized sequences of Sender reports, types and Receiver signals respectively by $(\tilde{\theta})^{\mathcal{G}(j)+T-1|\mathcal{G}(j)} = \left(\tilde{\theta}_{\mathcal{G}(j)}, \tilde{\theta}_{\mathcal{G}(j)+1}, \ldots, \tilde{\theta}_{\mathcal{G}(j)+T-1}\right)$, $(\theta)^{\mathcal{G}(j)+T-1|\mathcal{G}(j)} = \left(\theta_{\mathcal{G}(j)}, \theta_{\mathcal{G}(j)+1}, \ldots, \theta_{\mathcal{G}(j)+T-1}\right)$ and $(\omega)^{\mathcal{G}(j)+T-1|\mathcal{G}(j)} = \left(\omega_{\mathcal{G}(j)}, \omega_{\mathcal{G}(j)+1}, \ldots, \omega_{\mathcal{G}(j)+T-1}\right)$. Similarly, in punishment phases $\mathcal{B}_j$ define $(\tilde{\theta})^{\mathcal{G}(j)+(1+\beta)T|\mathcal{G}(j)+T}$, $(\theta)^{\mathcal{G}(j)+(1+\beta)T|\mathcal{G}(j)+T}$, $(\omega)^{\mathcal{G}(j)+(1+\beta)T|\mathcal{G}(j)+T}$. Due to the recursive nature of the equilibrium we establish below, the same arguments will hold across all $\mathcal{G}_j$, $\mathcal{B}_j$, $j = 1, 2, \ldots$, 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 $\left(\theta_t^i, \omega_t^j\right)$ given sequences $(\theta)_{\mathcal{G}}^{T-1}$ and $(\omega)_{\mathcal{G}}^{T-1}$ as $\mathcal{F}_T\left(\theta^i, \omega^j\right) := \frac{\sum_{t=0}^{T-1} \mathbf{1}\left(\theta_t^i, \omega_t^j\right)}{T}$. Similarly, given $\left(\tilde{\theta}\right)_{\mathcal{G}}^{T-1}$, $(\omega)_{\mathcal{G}}^{T-1}$, let the empirical frequency of joint observation $\left(\tilde{\theta}_t = \theta^i, \omega_t^j\right)$ be $\tilde{\mathcal{F}}_T\left(\tilde{\theta} = \theta^i, \omega^j\right) := \frac{\sum_{t=0}^{T-1} \mathbf{1}\left(\tilde{\theta}_t = \theta^i, \omega_t^j\right)}{T}$. Finally, given sequences $\left(\tilde{\theta}\right)_{\mathcal{G}}^{T-1}$, $(\theta)_{\mathcal{G}}^{T-1}$, let $b\left(\tilde{\theta} = \theta^i \mid \theta^k\right)$, $i, k \in \{1, 2, \ldots, N\}$, be the frequency of reports $\tilde{\theta} = \theta^i$ when S observes θ^k,

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

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 Sender adopts a truth-telling strategy within review period $\mathcal{G}_0$, he passes the review with probability at least $\Pr\left(\cap_{i,j \in \{1,2,\ldots,N\}} \left\{\left|\tilde{\mathcal{F}}_T\left(\tilde{\theta} = \theta^i, \omega^j\right) - f\left(\theta^i, \omega^j\right)\right| \leq \chi\right\}\right) \geq 1 - \epsilon_{\overline{T}}$; *(ii)* $\Pr\left(\cap_{i,j \in \{1,2,\ldots,N\}} \left\{\left|\mathcal{F}_T\left(\theta^i\right) - \mu_0^i\right| \leq \chi\right\}\right) \geq 1 - \epsilon_{\overline{T}}$; *(iii)* for any sequences $(\tilde{\theta})_{\mathcal{G}}^{T-1}$, $(\theta)_{\mathcal{G}}^{T-1}$ such that $\left|\mathcal{F}_T\left(\theta^i\right) - \mu_0^i\right| \leq \chi$, $i = 1, 2, \ldots, N$, if $\left|\left(b\left(\tilde{\theta} = \theta^k \mid \theta^i\right)\right)_{k=1}^{N} - \mathbf{e}_{\theta^i}\right| > \xi_T$, then S fails the review with probability at least $1 - \epsilon_{\overline{T}}$, where $\mathbf{e}_{\theta^i} = (0, 0 \ldots, 1, \ldots, 0)$ 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 $\hat{\theta}_t\left((\tilde{\theta})^{t-1}, (\theta)^{t-1}, \theta_t\right) = \theta_t$. Under the truthful strategy, $\left(\tilde{\theta}, \omega\right)$ is a sequence of independent Bernoulli trials. By Chebyshev's inequality,

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

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

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

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

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

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

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

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

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

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

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

$$\Pr\left(\left|\tilde{\mathcal{F}}_{T'}\left(\tilde{\theta}=\theta^i,\omega^j\right) - \mathbb{E}\left[\tilde{\mathcal{F}}_{T'}\left(\tilde{\theta}=\theta^i,\omega^j\right) \mid \left(\tilde{\theta},\theta\right)_{\mathcal{G}}^{T-1}\right]\right| \le \chi\left(T'\right) \mid \left(\tilde{\theta},\theta\right)_{\mathcal{G}}^{T-1}\right) \ge 1 - \frac{\psi^{-2}}{4\left(T'\right)^{1-2y}} \tag{13}$$

Now, since $\left|\mathcal{F}_{T'}\left(\theta^k\right)-\mu_0^k\right| \le \chi\left(T'\right)$, and by definition $0 \le \sum_{\theta^k\in\Theta} b\left(\tilde{\theta}=\theta^i \mid \theta^k\right) f\left(\omega^j \mid \theta^k\right) < \sum_{\theta^k\in\Theta} b\left(\tilde{\theta}=\theta^i \mid \theta^k\right) \le N$, we can bound

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

which follows after recalling that $f\left(\theta^k,\omega^j\right)=\mu_0^i f\left(\omega^j \mid \theta^k\right)$, $\forall\theta^k$. Using the triangle inequality and (13), we can bound $\Pr\left(\tilde{A}_{i,j} \mid \left(\tilde{\theta},\theta\right)_{\mathcal{G}}^{T'-1}\right) \ge 1-\frac{\psi^{-2}}{4(T')^{1-2y}}$, where

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

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

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

$$\boldsymbol{F}_{\theta,\omega}\cdot\boldsymbol{b}_i - (N+1)\,\chi\left(T'\right)\cdot\mathbf{1} \le \ \tilde{\boldsymbol{\mathcal{F}}}_{\boldsymbol{T'}} \ \le \boldsymbol{F}_{\theta,\omega}\cdot\boldsymbol{b}_i + (N+1)\,\chi\left(T'\right)\cdot\mathbf{1} \tag{14}$$

where $\boldsymbol{F}_{\theta,\omega} = \begin{pmatrix} f(\theta^1,\omega^1) & \cdots & f\left(\theta^N,\omega^1\right) \\ \vdots & \ddots & \vdots \\ f\left(\theta^N,\omega^1\right) & \cdots & f\left(\theta^N,\omega^N\right) \end{pmatrix}$, $\tilde{\boldsymbol{\mathcal{F}}}_{\boldsymbol{T'}} = \begin{pmatrix} \tilde{\mathcal{F}}_{T'}(\theta^i,\omega^1) \\ \vdots \\ \tilde{\mathcal{F}}_{T'}\left(\theta^i,\omega^N\right) \end{pmatrix}$, $\boldsymbol{b}_i = \begin{pmatrix} b(\theta^i \mid \theta^1) \\ \vdots \\ b\left(\theta^i \mid \theta^N\right) \end{pmatrix}$ and $\mathbf{1} = \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix}$.

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

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

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

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

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

Selecting $\hat{T} = \max\left\{\overline{T}, T_\xi\right\}$, $\hat{\chi} = \chi\left(\hat{T}\right)$ establishes claim *(iii)*.

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

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

where $\hat{\mu} \in \Delta\Theta$.[40] Let Q_t denote the set of all pairs of sequences $q^t = \left((\tilde{\theta})^{t-1}, (\theta)^t\right)$, and $Q = \cup_{t\in\{0,1,\ldots,T-1\}} Q_t$. With some notational abuse, let $\tilde{\Theta}^Q$ denote the set of feasible reporting *plans* within phase $\mathcal{G}$, $\tilde{\theta}^Q := \left(\tilde{\theta}_t : \Theta^{t-1} \times \Theta^t \to \Theta\right)_{t=0}^{T-1}$. Given a reporting plan $\tilde{\theta}_t(q^t)$ and report function σ, we write the induced lottery over $\{\mu_1, \ldots, \mu_N\}$, conditional on $\left(\tilde{\theta}\right)^{t-1}$, $(\theta)^{t-1}$ as $\lambda_i(q^t) := \Pr\left(\tilde{\mu} = \mu_i \mid \tilde{\theta}^{t-1}, \theta^{t-1}\right)$. Associated with this lottery, define the induced lottery over Receiver posteriors as $\hat{\lambda}$, where $\hat{\lambda}_i := \Pr\left(\hat{\mu} = \hat{\mu}_i \mid \tilde{\theta}^{t-1}, \theta^{t-1}\right)$. Note that in any PBE, $\hat{\lambda} \in \Lambda(\mu_0)$.

Clearly, in any $\mathcal{B}$ phase, truthful reporting by 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 $\sum_{i=1}^N \underline{\lambda}_j v\left(\underline{\mu}_j\right) = 0$.

Consider 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 Sender's payoffs at the beginning of phase $\mathcal{G}$ recursively as

$$\mathcal{V}_{\mathcal{G}} = \max_{\tilde{\theta}^Q \in \Theta^Q} \sum_{t=0}^{T-1} \delta^t \mathbb{E}\left[\sum_{i=1}^N \hat{\lambda}_i\left(\tilde{\theta}^t\left(q^t\right)\right) v\left(\hat{\mu}_{t,i}\right)\right] + \delta^T\left[(1-\tilde{\varphi}) + \tilde{\varphi}\delta^{\beta T}\right] \mathcal{V}_{\mathcal{G}} \tag{15}$$

(15) is a standard dynamic programming problem with a finite set of states - therefore, a stationary optimal strategy for Sender exists. Letting the best response function be $\tilde{\theta}^t = \overline{\theta}^t$, we can equivalently write (15) as[41]

$$(1-\delta)\mathcal{V}_{\mathcal{G}} = (1-\delta)\frac{\sum_{t=0}^{T-1} \delta^t \mathbb{E}\left[\sum_{i=1}^N \hat{\lambda}_i\left(\overline{\theta}^t\left(q^t\right)\right) v\left(\hat{\mu}_{t,i}\right)\right]}{1 - (1-\tilde{\varphi})\delta^T - \tilde{\varphi}\delta^{(1+\beta)T}} \tag{16}$$

or

$$(1-\delta)\mathcal{V}_{\mathcal{G}} = \frac{(1-\delta)}{1-\delta^T}\left[\sum_{t=0}^{T-1} \delta^t \mathbb{E}\left[\sum_{i=1}^N \hat{\lambda}_i\left(\overline{\theta}^t\left(q^t\right)\right) v\left(\hat{\mu}_{t,i}\right)\right] - \delta^T\left(1-\delta^{\beta T}\right)\tilde{\varphi}(1-\delta)\mathcal{V}_{\mathcal{G}}\right] \tag{17}$$

We now argue that, for any $\underline{\varphi} \in (0,1)$ and $\varepsilon > 0$ and $\eta > 0$, there exist $\overline{\overline{\chi}}$, $\overline{\overline{T}}$, r, $\overline{\beta}$, $\delta_{\overline{\overline{T}}}$, such that if Receiver posteriors given Mediator message $\tilde{\mu}_i$ satisfy $|\hat{\mu}_i - \tilde{\mu}_i| < \eta$ for all $i \in \{1, 2, \ldots, N\}$, then *(i)* 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 $\delta \geq \delta_{\overline{\overline{T}}}$.

Fix a Mediator report function $r\left(\mathcal{G}, \tilde{\theta}_t\right) = \sigma' : \Theta \to \Delta\Theta$, with support $\{\mu_1', \ldots, \mu_N'\}$ and

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

[41] We suppress explicit dependence of $\overline{\theta}$ on $\hat{\chi}$, $\hat{T}$, r, $\hat{\mu}$, δ, β for notational ease.

induced lottery λ', where λ' satisfies $\lambda' \in X(\lambda^{\star\star})$ for some optimal experiment under commitment, $\lambda^{\star\star} \in \Lambda(\mu_0)$, 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 3. Suppose further that Receiver posterior beliefs, $\hat{\mu}'(\tilde{\mu})$, satisfy $|\hat{\mu}_i' - \tilde{\mu}_i| < \eta'$, for all $i \in \{1, 2, \ldots, N\}$, where η' is small 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 $\hat{\mu}_i' \in N_{\eta'}(\tilde{\mu}_i)$, $i = 1, \ldots, N$.

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 (17) as $\delta \to 1$ 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 Sender's payoff as

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

where $\overline{v} := \max\{v(\mu_i)\}_{i=1}^{N}$. Clearly, given a $\overline{\varphi} > 0$, we can find a $\overline{\beta} \in \mathbb{N}_{>0}$ such that $\overline{v} - \overline{\beta}\overline{\varphi}\left(\hat{v}(\mu_0) - \frac{3\varepsilon}{4}\right) \leq \hat{v}(\mu_0) - \frac{3\varepsilon}{4}$ - a contradiction to the lower bound, $(1-\delta)\mathcal{V}_{\mathcal{G}} \geq \hat{v}(\mu_0) - \frac{3\varepsilon}{4}$. Thus, for any $\beta \geq \overline{\beta}$, it cannot be optimal (in the limit) for Sender to use a reporting plan $\tilde{\theta}_t(q^t)$ in phase $\mathcal{G}$ that $\tilde{\varphi} > \overline{\varphi}$. We now establish that such a lower bound can be imposed, given $\overline{\beta}$. Fixing $\overline{\beta} > 0$, note that we can bound Sender's limit payoff, $\lim_{\delta \to 1}(1-\delta)\mathcal{V}_{\mathcal{G}}$, as $\delta \to 1$ in each $\mathcal{G}$-phase by Sender's limit payoff from truthful reporting. From equation (16),

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

By choosing $\overline{\overline{T}}$ sufficiently high, and $\overline{\overline{\chi}} = \chi\left(\overline{\overline{T}}\right)$ we can set $\epsilon_{\overline{\overline{T}}}$ small enough that $\frac{\hat{v}(\mu_0) - \frac{\varepsilon}{2}}{1 + \epsilon_T \beta} \geq \hat{v}(\mu_0) - \frac{3\varepsilon}{4}$. 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 are are only finitely many such choices of recursive strategy, because each phases has finitely many nodes and messages. Given $\overline{\overline{T}}$, we can therefore select $\delta_{\overline{\overline{T}}} < 1$ 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 $(1-\delta)\mathcal{V}_{\mathcal{G}} \geq v(\hat{\mu}_0) - \varepsilon$, for all $\delta \geq \delta_{\overline{\overline{T}}}$ - as required.

Given that A is countable, Assumption 3 ensures there exists an $\eta'' > 0$ such that $a^\star(\hat{\mu}_i') = a^\star(\tilde{\mu}_i)$, 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 $\overline{\varphi}$ such that Receiver beliefs are guaranteed to satisfy $|\hat{\mu}_i' - \tilde{\mu}_i| < \eta''$ in equilibrium - Sender's and Receivers' optimal strategies are then mutual best responses, and beliefs are correct.

For any $\overline{\varphi} \in (0,1)$, we have shown there exist $\overline{\overline{\chi}}$, $\overline{\overline{T}}$, r, $\overline{\beta}$, $\delta_{\overline{\overline{T}}}$ such that $\tilde{\varphi} \leq \underline{\varphi}$, and *(ii)* $(1-\delta)\mathcal{V}_{\mathcal{G}} \geq \hat{v}(\mu_0) - \varepsilon$, for all $\delta \geq \delta_{\overline{\overline{T}}}$. Additionally, for any $\epsilon_{T^\star} \leq \epsilon_{\overline{\overline{T}}}$, $\xi_{T^\star} > 0$ we can choose $\chi(T^\star)$, $T^\star$, where $T^\star = \max\left\{\overline{\overline{T}}, \hat{T}\right\}$, to ensure that sequences $(\theta)_{\mathcal{G}}^{T-1}$ satisfying $\left|\mathcal{F}_T(\theta^i) - \mu_0^i\right| \leq \chi(T^\star)$, for all $i = 1, 2, \ldots, N$, occur with probability greater than $1 - \epsilon_{T^\star}$ in review period $\mathcal{G}$ (claim *(i)*); conditional on any such sequence, a reporting strategy that involves $|\boldsymbol{b}_i - \mathbf{e}_{\theta^i}| > \xi_T$ fails the review with probability at least $1 - \epsilon_{\overline{T}}$ (claim *(iii)*).

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

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

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

$$\Pr\left(\cap_i \left\{\left|\mathcal{F}_{T^\star}(\theta) - \mu_0^i\right| \leq \chi^\star\right\}, |\boldsymbol{b}_i - \boldsymbol{e}_{\theta_i}| > \xi_{T^\star}\right)(1 - \epsilon_{T^\star}) \leq \Pr\left(S \text{ fails review at } T^\star - 1\right) \leq \overline{\varphi}$$

and

$$\Pr\left(\cup_i \left\{\left|\mathcal{F}_{T^\star}(\theta) - \mu_0^i\right| > \chi^\star\right\}, |\boldsymbol{b}_i - \boldsymbol{e}_{\theta_i}| > \xi_{T^\star}\right) \leq \Pr\left(\cup_i \left|\mathcal{F}_{T^\star}(\theta) - \mu_0^i\right| > \chi^\star\right) \leq \epsilon_{T^\star}$$

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

$$T_L(\xi) := \max_{\mathcal{T}\in 2^{\{0,1,\ldots,T^\star-1\}}} |\mathcal{T}| \tag{18}$$

s.t.

$$\left|\boldsymbol{b}_i^{\mathcal{T}} - \boldsymbol{e}_{\theta_i}\right| \geq \xi^{\frac{1}{2}}$$

where $|\mathcal{T}|$ denotes the cardinality of set $\mathcal{T}$. Since problem (18) 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 $|\boldsymbol{b}_i - \boldsymbol{e}_{\theta_i}| \leq \xi$, 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 $(|\boldsymbol{B}-\boldsymbol{E}|\leq\xi):=\left\{\left(\tilde{\theta}\left(q^{t}\right),\theta\right)_{\mathcal{G}}^{T-1}:|\boldsymbol{b}_{i}-\boldsymbol{e}_{\theta_{i}}|\leq\xi,\forall i\in\{1,\ldots,N\}\right\}$, we can write posterior beliefs given message $\tilde{\mu}$ as

$$\begin{aligned}
\Pr\left(\theta_{\tilde{\pi}(t)}=\theta_{i}\mid\mathcal{G},\tilde{\mu}\right) &= \Pr\left(\theta_{\tilde{\pi}(t)}=\theta_{i},|\boldsymbol{B}-\boldsymbol{E}|\leq\xi^{\frac{1}{2}},\tilde{\pi}\left(t\right)\notin\mathcal{T}_{L}^{\xi}\mid\mathcal{G},\tilde{\mu}\right) \\
&+ \Pr\left(\theta_{\tilde{\pi}(t)}=\theta_{i},|\boldsymbol{B}-\boldsymbol{E}|\leq\xi^{\frac{1}{2}},\tilde{\pi}\left(t\right)\in\mathcal{T}_{L}^{\xi}\mid\mathcal{G},\tilde{\mu}\right) \\
&+ \Pr\left(\theta_{\tilde{\pi}(t)}=\theta_{i},|\boldsymbol{B}-\boldsymbol{E}|>\xi^{\frac{1}{2}}\mid\mathcal{G},\tilde{\mu}\right)
\end{aligned}\tag{19}$$

The final two terms can be bounded respectively by

$$\begin{aligned}
\Pr\left(\theta_{\tilde{\pi}(t)}=\theta_{i},|\boldsymbol{B}-\boldsymbol{E}|\leq\xi^{\frac{1}{2}},\tilde{\pi}\left(t\right)\in\mathcal{T}_{L}^{\xi}\mid\mathcal{G},\tilde{\mu}\right) &\leq \Pr\left(|\boldsymbol{B}-\boldsymbol{E}|\leq\xi^{\frac{1}{2}},\tilde{\pi}\left(t\right)\in\mathcal{T}_{L}^{\xi}\mid\mathcal{G},\tilde{\mu}\right) \\
&\leq \Pr\left(\tilde{\pi}\left(t\right)\in\mathcal{T}_{L}^{\xi}\mid|\boldsymbol{B}-\boldsymbol{E}|\leq\xi^{\frac{1}{2}},\mathcal{G},\tilde{\mu}\right)
\end{aligned}$$

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

As ξ, $\epsilon_T\to 0$, we have shown above that[42]

$$\Pr\left(\tilde{\pi}\left(t\right)\in\mathcal{T}_{L}^{\xi}\mid|\boldsymbol{B}-\boldsymbol{E}|\leq\xi^{\frac{1}{2}},\mathcal{G}\right),\Pr\left(|\boldsymbol{B}-\boldsymbol{E}|>\xi^{\frac{1}{2}}\mid\mathcal{G}\right)\to 0.$$

Further it is straightforward to show that $\Pr\left(\tilde{\mu}_{\tilde{\pi}(t)}\mid\mathcal{G}\right)\to\tilde{\lambda}\left(\tilde{\mu}\left(\theta_{t}\left(q^{t}\right)\right)\right)$ as ξ, $\epsilon_T\to 0$, where $\theta_t\left(q^t\right)$ is the truthful reporting strategy, and $\tilde{\lambda}\left(\tilde{\mu}\right)$ is the induced frequency of message $\tilde{\mu}$ when Mediator uses message rule $\tilde{\sigma}\left(\tilde{\theta},\mathcal{G}\right)$. Since there are only finitely many messages sent in $\tilde{\sigma}$ (Lemma **XXXX**), it is without loss that $\tilde{\lambda}\left(\tilde{\mu}\right)>0$ for any $\tilde{\mu}\in supp\left(\tilde{\sigma}\right)$. Thus, using Bayes' rule, $\Pr\left(|\boldsymbol{B}-\boldsymbol{E}|>\xi^{\frac{1}{2}}\mid\mathcal{G},\tilde{\mu}\right)=\frac{\Pr\left(|\boldsymbol{B}-\boldsymbol{E}|>\xi^{\frac{1}{2}}|\mathcal{G}\right)\Pr\left(\tilde{\mu}||\boldsymbol{B}-\boldsymbol{E}|>\xi^{\frac{1}{2}},\mathcal{G}\right)}{\Pr(\tilde{\mu}|\mathcal{G})}\leq\frac{\Pr\left(|\boldsymbol{B}-\boldsymbol{E}|>\xi^{\frac{1}{2}}|\mathcal{G}\right)}{\Pr(\tilde{\mu}|\mathcal{G})}\to 0$. Similarly, $\Pr\left(\tilde{\pi}\left(t\right)\in\mathcal{T}_{L}^{\xi}\mid|\boldsymbol{B}-\boldsymbol{E}|\leq\xi^{\frac{1}{2}},\mathcal{G},\tilde{\mu}\right)\to 0$. Thus, as ξ, $\epsilon_T\to 0$, (19) implies that

$$\Pr\left(\theta_{\tilde{\pi}(t)}=\theta_{i}\mid\mathcal{G},\tilde{\mu}\right)\to\Pr\left(\theta_{\tilde{\pi}(t)}=\theta_{i}\mid\mathcal{G},\tilde{\mu},|\boldsymbol{B}-\boldsymbol{E}|\leq\xi^{\frac{1}{2}},\tilde{\pi}\left(t\right)\notin\mathcal{T}_{L}^{\xi}\right)\Pr\left(|\boldsymbol{B}-\boldsymbol{E}|\leq\xi^{\frac{1}{2}},\tilde{\pi}\left(t\right)\notin\mathcal{T}_{L}^{\xi}\mid\mathcal{G},\tilde{\mu}\right)$$

But for any time period $\tilde{\pi}\left(t\right)$ in which $\Pr\left(\tilde{\mu}_{\pi(t)}=\tilde{\mu}\right)>0$, $\Pr\left(|\boldsymbol{B}-\boldsymbol{E}|\leq\xi^{\frac{1}{2}},\tilde{\pi}\left(t\right)\notin\mathcal{T}_{L}^{\xi}\mid\mathcal{G},\tilde{\mu},\tilde{\pi}\left(t\right)\right)\to 1$ since otherwise, $1-\Pr\left(|\boldsymbol{B}-\boldsymbol{E}|\leq\xi^{\frac{1}{2}},\tilde{\pi}\left(t\right)\notin\mathcal{T}_{L}^{\xi}\mid\mathcal{G},\tilde{\mu}\right)$ would be bounded below by some $\Pr\left(\tilde{\mu}_{\pi(t)}=\tilde{\mu}\right)x$, $x>0$ - a contradiction to $\lim_{\xi,\epsilon_T\to 0}\Pr\left(|\boldsymbol{B}-\boldsymbol{E}|\leq\xi^{\frac{1}{2}},\tilde{\pi}\left(t\right)\notin\mathcal{T}_{L}^{\xi}\mid\mathcal{G},\tilde{\mu}\right)=$

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

1. Finally,by Bayes' Rule,

$$\Pr\left(\theta_{\tilde{\pi}(t)} = \theta_i \mid \mathcal{G}, \tilde{\mu}, |\boldsymbol{B} - \boldsymbol{E}| \leq \xi^{\frac{1}{2}}, \tilde{\pi}(t) \notin \mathcal{T}_L^{\xi}\right) = \frac{\Pr\left(\theta_{\tilde{\pi}(t)} = \theta_i \mid \mathcal{G}, |\boldsymbol{B} - \boldsymbol{E}| \leq \xi^{\frac{1}{2}}, \tilde{\pi}(t) \notin \mathcal{T}_L^{\xi}\right) \Pr\left(\tilde{\mu} \mid \theta_{\tilde{\pi}(t)}\right.}{\Pr\left(\tilde{\mu} \mid \mathcal{G}, |\boldsymbol{B} - \boldsymbol{E}| \leq \xi^{\frac{1}{2}}, \tilde{\pi}(t)\right.}$$

Using an identical argument to that above, $\Pr\left(\theta_{\tilde{\pi}(t)} = \theta_i \mid \mathcal{G}, |\boldsymbol{B} - \boldsymbol{E}| \leq \xi^{\frac{1}{2}}, \tilde{\pi}(t) \in \mathcal{T}_L^{\xi}\right) \to \Pr\left(\theta_{\tilde{\pi}(t)} = \theta_i\right)$. Furthermore, for any θ_i, $\xi \to 0$ implies $\boldsymbol{b}_i^{\hat{\mathcal{T}}} \to \boldsymbol{e}_{\theta_i}$, for $\hat{\mathcal{T}} = \{0, 1, \ldots, T-1\} / \mathcal{T}_L^{\xi}$. By continuity,

$$\lim_{\xi, \epsilon_T \to 0} \Pr\left(\tilde{\mu} \mid \theta_{\tilde{\pi}(t)} = \theta_i, \mathcal{G}, |\boldsymbol{B} - \boldsymbol{E}| \leq \xi^{\frac{1}{2}}, \tilde{\pi}(t) \notin \mathcal{T}_L^{\xi}\right) = \tilde{\sigma}(\tilde{\mu} \mid \theta_i) \tag{20}$$

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

Finally, (20) implies that we can find ξ, ϵ_T small enough that $|\hat{\mu} - \tilde{\mu}| < \eta''$, for any $\eta'' > 0$. Selecting $T^\star$, $\chi(T^\star)$ large enough ensures such ξ, ϵ_T can be found. This can be sustained as shown above by a bound on the discount rate, $\delta \geq \delta_{T^\star}$ for some $\delta_{T^\star} < 1$.

$\square$

Appendix B: Definition of Review Mechanisms

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

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

$$\mathcal{B}(j) \in \{0, \beta T\} \tag{21}$$

$$\mathcal{G}(j+1) = \mathcal{G}(j) + T + \mathcal{B}(n) \tag{22}$$

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

$$\mathcal{G}(1) = 0 \tag{23}$$

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

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

If $\mathcal{B}(j) > 0$, then we say that review phase j is followed by an associated *punishment phase* $\mathcal{B}^j := \{\mathcal{G}(j) + T, \mathcal{G}(j) + T + 1, \ldots, \mathcal{G}(j) + T + \mathcal{B}(j)\}$; 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 $(\mathcal{G}_j, \mathcal{B}_j)_{j=1}^{\iota}$ satisfying (21)-(23), $\iota = 1, 2, \ldots$, and $t \in \mathcal{G}_\iota \cup \mathcal{B}_\iota$ be $\mathcal{A}_t$. Finally, let $\pi_j : \mathcal{G}^j \to \mathcal{G}^j$ 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 $\Pr(\pi_j = \tilde{\pi}) = \frac{1}{T!}$. Given $\tilde{\pi}$, the Mediator commits to permute the ordering of Receivers $R_{\mathcal{G}(j)}, \ldots, R_{\mathcal{G}(j)+T-1}$ 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.[43] Similarly, the Mediator uniformly permutes the ordering of Receivers in punishment phases.

[43]It 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.

At each time $t \in \{0, 1, \dots\}$, the review mechanism asks 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, $\left(\tilde{\theta}_\tau\right)^{t-1} = \left(\tilde{\theta}_1, \dots, \tilde{\theta}_{t-1}\right)$, and the history of phases announced by the Mediator, $(\mathcal{G}_j, \mathcal{B}_j)_{j=1}^{\iota}$ 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 $r : \mathcal{G}^\iota \cup \mathcal{B}^\iota \times \Theta \to \Delta M$, 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 $\tilde{\theta}_t : \Theta^t \times \Theta^{t-1} \times \mathcal{A}_{t-1} \times \mathcal{P}_t \to \Delta\Theta$, 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 $\rho : \mathcal{P}_t \times M \to A$, 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 $(\omega_\tau)^{\mathcal{G}(j)+T-1} = \left(\omega_1, \omega_2, \dots, \omega_{\mathcal{G}(j)+T-1}\right)$, reports $\left(\tilde{\theta}_\tau\right)^{\mathcal{G}(j)+T-1}$ 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 $\kappa\left(\omega^k, (\omega_\tau)^t, \mathcal{G}^j\right) := \left\{\tau : \tau \in \mathcal{G}^j \ s.t.\ \omega_t = \omega^k\right\}$, and the corresponding empirical frequency of reports, $\tilde{\theta}^i \in \Theta$, as:

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

With some notational abuse, we will often suppress the dependence of $\mathcal{F}$ on $\left(\tilde{\theta}\right)^{\mathcal{G}(j)+T-1}$, $\mathcal{G}^j$ and simply write $\mathcal{F}\left(\tilde{\theta}^i \mid \omega^k\right)$ where clear. Given $\mathcal{F}$, we can calculate, for each $\theta^i \in \Theta$, $\omega^k \in \Omega$, the difference $d\left(\theta^i, \omega^k\right)$ between the empirical frequency and theoretical probability of observing θ^i given ω^k, as

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

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 $d\left(\theta^i, \omega^k\right)$

for all $\theta^i \in \Theta$, and all $\omega^k \in \Omega$. The Sender *passes the review phase* $\mathcal{G}^j$ if

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

for all $\theta^i \in \Theta$, and all $\omega^k \in \Omega$ satisfying $\left|\kappa\left(\omega^k, \mathcal{G}^j\right)\right| > \underline{\kappa}$, 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, $\mathcal{B}(j) = \beta T$.