Honestly Dishonest: A Solution to the Commitment Problem in Bayesian Persuasion*

James Best † Daniel Quigley‡

This version: December 9, 2016

Abstract

Optimal persuasion, *à la* Kamenica-Gentzkow (2011), require Senders to commit to reporting strategies; one potential source of such commitment is repeated interaction. We study a model in which a long lived Sender plays a cheap talk game with a sequence of short lived Receivers who observe the joint histories of reports and states. In this setting, we show we can achieve optimal persuasion if and only if honest reporting is optimal. However, as demonstrated in the persuasion literature, honest reporting is generally not optimal for the Sender. We then show how optimal persuasion can always be retrieved by altering the game so it has a property that we call "honestly dishonest". We show that we can make a game honestly dishonest by using cryptographic technologies or mediators. We then give several examples of games that are honestly dishonest in the first place.

1 Introduction

Being able to believe and act on the claims of others is a corner stone of a well-functioning society. A major source of such trust is that people want to be trusted in the future; the cost of lying today is that they may not be believed tomorrow. Given this, should people always tell the truth or can they gain from being systematically dishonest?

*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)

We develop a general model which allows us to examine this question. Importantly, we do so without the presence of the commitment assumption in Kamenica Gentzkow (2011). A patient Sender plays a cheap talk game with a sequence of short lived Receivers. Each period a state that is payoff relevant to the Receiver is realised according to an i.i.d. process. This state is observed by the Sender; he then sends a message about the state to the Receiver; she then takes an action based on her belief about the state; finally, the state is revealed. The Sender's stage payoff is the result of the action taken by the Receiver. The Receiver's belief about the state results from the message she receives and her beliefs about the Sender's strategy. Each new Receiver observes the joint history of messages and states which she uses to form her beliefs about the Sender's strategy. The Sender then, must look further than today's Receiver and also consider how his advice will affect 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.

As a preliminary result we show the threat of permanently ignoring the Sender's messages can sustain honest equilibria for sufficiently patient Senders. Where an honest equilibria is one in which the Sender sends one, and only one, message in any particular state.[1] Next we show that, *à la* Kamenica Gentzkow (2011), we can also sustain equilibria where the Sender can commit to any reporting policy at *some* stage in the game. However, when that reporting policy is dishonest this implies that in some states of the world they are mixing between the truth and a lie - where the lie is the more profitable message to send at the stage game. To be able to commit to such a policy the Sender has to be indifferent between telling the truth and lying - this indifference is maintained by on path punishment whenever the Sender lies. Such dishonest equilibria are then characterised by trust cycles as Receivers punish the Sender by ignoring him for some time after a lie.[2]

The existence of these on path punishments lead us to our next main result: the Sender can only achieve the optimal average payoff if the optimal strategy in the commitment game is honest. Note, most of the literature examines cases where the optimal strategy is not honest. Further, we show when repetition is the only commitment mechanism for the Sender that honest equilibria Pareto dominate dishonest equilibria in many of the canonical cases. Unfortunately, in many of these cases, the Sender has little incentive to move to an honest equilibrium. This should provide a note of caution about encouraging attempts at persuasion in environments where repeated cheap talk is the form of commitment.

[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.

[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.

However, our third main result gives us reason for hope: we can retrieve the payoffs available under full commitment if a game has the property of being "honestly dishonest." To describe this property we will consider how we can make the game honestly dishonest with what we call a "Coin and Cup" (CnC). A CnC is a payoff irrelevant random variable which realizes some value each period; the Sender observes the CnC before giving advice and the Receiver does not; the value of the CnC is publicly revealed at the same time as the payoff relevant state. Receiver's now observe the joint history of advice, states and realizations of the CnC. The Sender can condition pure strategies on the CnC, for any realization of the state, that are ex-ante equivalent to mixing from the perspective of the Receiver. However, ex-post, Receivers can see if he deviated from his pure strategy and only punish him if he lies for disallowed values of the CnC. That is, they can monitor whether he is being honestly dishonest and so never need to punish him on the equilibrium path. It is worth noting here that the field of cryptography already uses a technology analagous to a CnC for other purposes; it is called a 'commitment scheme' and can be easily placed in a blockchain to generate such a publicly verifiable record.

A game is honestly dishonest then, if the Sender can adopt a reporting strategy of sending multiple messages in some payoff relevant state of the world according to a rule that is stochastic from the Receiver's perspective ex-ante but is deterministic and verifiable after the Receiver has acted. We will show later that games can also be honestly dishonest when there is a mediator or many Receivers acting simultaneously. Finally, we will examine a case where institutional changes have been pushed by advisers within finance that allow their customers to verify the rate at which they have "lied". The effect of this change nicely illustrates the results in this paper.

This paper is, in part, an analysis of how the commitment assumption layed out in Kamenica and Gentzkow [2011] and applied in the subsequent literature (Rayo and Segal, 2010; Perez-Richet, 2014; Kolotilin, 2015; Taneva, 2015; Tamura, 2016) can be microfounded by embedding cheap talk within a dynamic setting. Dynamic persuasion has been analysed in Kremer et al. [2014], Ely [2015], Bizzotto et al. [2016], and elsewhere. However, unlike in this paper, these are analyses of optimal persuasion when there is commitment at the stage game.

The closest work to ours are Sobel [1985], Hermalin [2007], and Margaria and Smolin [2015]. Sobel (1985) examines a finitely repeated cheap talk where the Sender is either a 'friend' whose preferences are aligned with the Receiver or an 'enemy' who's preferences are opposed. The desire to appear to be a friend can cause even enemies to be truthful. We do not rely on multiple types of Senders in our paper. More importantly we examine a different question: when beliefs about the Sender's reporting strategy can discipline behaviour today what is the best the Senders and Receivers can achieve.

Hermalin (2007) examines a game where a Leader has private information about the value of a public good and communicates with a team of myopic Followers using cheap talk and costly signalling. We generalise one of the results in this paper about truthful communication to a generalised setting. Unlike us Hermalin does not examine how a leader may do better through dishonest communication nor how we can improve upon the equilibria of a standard cheap talk game.

Margaria and Smolin's paper does not seek to answer the problem of persuasion as the Sender. They develop a folk theorem for a repeated cheap talk game with a long lived Receiver showing the players can attain any individual rational average payoffs. This sharp difference from our result follows from the Sender's ability to punish the Receiver by withholding information in the future. The threat of this punishment means the sender can extract far more from the Receiver than in a Bayesian persuasion game with full commitment. Our model, by focusing on Receiver's whose actions are determined only by their beliefs about the state, does seek to examine the question of persuasion. Moreover, we can also look at how small changes to the game are able to allow for greater levels of persuasion and communication: shifting out the Pareto frontier of the game.

The paper is organised as follows: in section 2 we provide a simple example that illustrates the main results of the paper; in section 3 we describe the full model; in section 4 we analyse the equilibrium and give the main results; in section 5 we look at some real world examples of games that may be honestly dishonest; and in section 6 we conclude.

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 proposal the Firm always gets zero; the Mayor gets a payoff of one when the project is Good and minus one when the project is Bad. Hence, the Mayor will only okay the project if she believes the project will be Good with at least a fifty percent probability. The red 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 fdsfdas

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

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

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

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

Figure 1

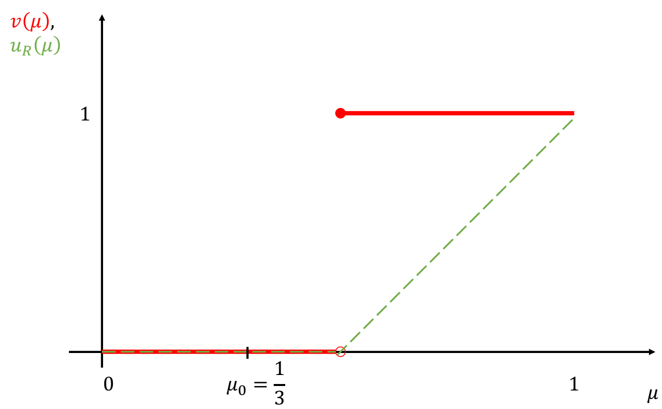

The Firm then sends a report of 'Accept' or 'Reject'. The firm commits to a policy of sending the 'Accept' report with some probability if the project is Good, and with a weakly lower probability when the project is Bad. The (Bayesian) Mayor, after any given report, forms posterior beliefs based on the 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. In which case the expected payoff to the Firm is just the probability of sending an accept, from (1) this is

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

Figure 2

(a)

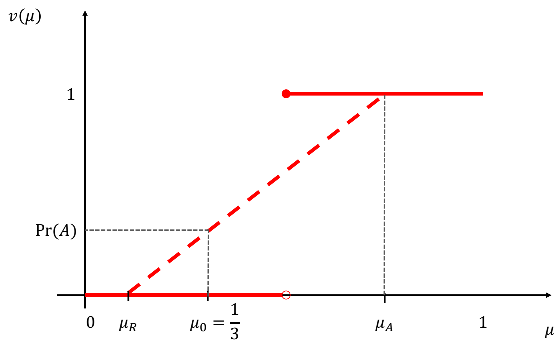

(b)

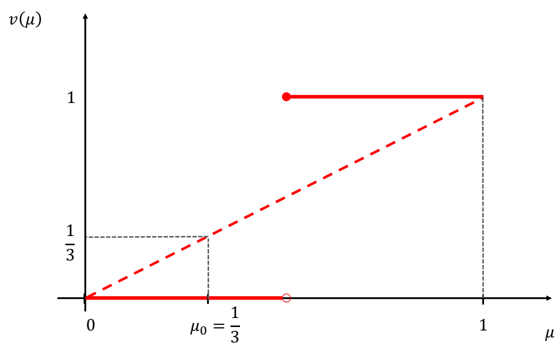

(c)

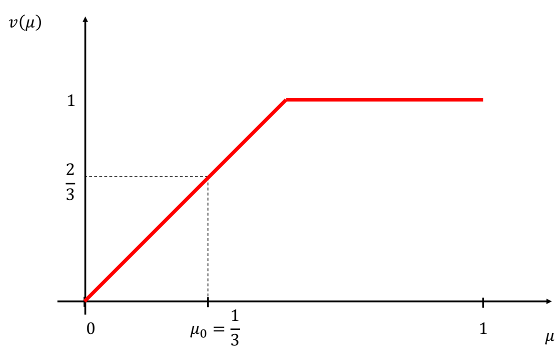

which is just the height of the line connecting the posteriors μ_R and μ_A evaluated at the prior μ_0 as shown in figure 2a below. From the Mayor's perspective, the best policy is Truth Telling -the Firm sends an Accept if and only if the project is Good, $(\mu_{R,}\mu_A) = (0, 1)$. This pays out for Good projects only and, 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's 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).

The ability to commit to a reporting policy in a one shot setting like that above may be infeasible. In such cases there is no equilibrium in which the Mayor will follow the advice of Accept. 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 whether it is trusted in the future there may be some 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, one Mayor per term. As before, the prior probability of the project being Good is $\mu_0 = 1/3$. Being short lived[3] each Mayor's decision rule and payoff will not differ from the full commitment case; a term t Mayor will accept the project if she believes it has at least a fifty percent chance of being Good: her posterior $\mu^t \geq 0.5$. The Mayor in each term t observes all the reports that the Firm has sent, the outcomes of all accepted projects, and she forms her beliefs about the Firms reporting 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

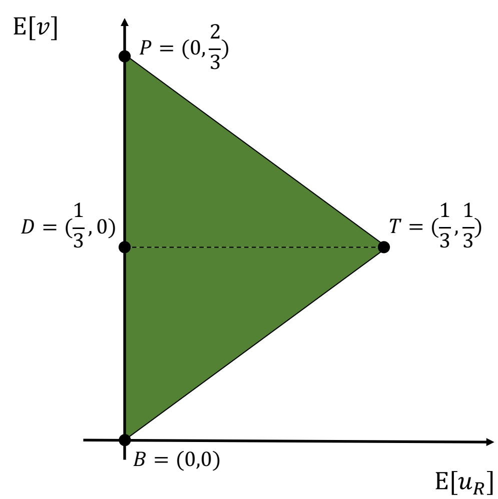

We will compare what can be achieved in this repeated setting to the baseline payoffs of full commitment described in Figure 3. B is the payoff from Babbling; T is the payoff from Truth Telling; and P is the payoff from optimal Persuasion, Firm reports Accept half of the time when the project is Bad and always when Good. 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, Mayor's 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 Mayor's 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:

[3]We restrict attention in the main body of this paper to one period Receivers because we want the Receivers' actions to be determined only by the Sender's ability to induce beliefs and not by the effect of Receivers' own actions on their continuation payoffs. Need to say this better.

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

If at some term t, the Firm learns the project is Bad but reports an Accept the Firm gains a payoff of one but loses the Truth Telling continuation payoff as the game moves to a Babbling equilibrium in which the Firm gets nothing. Hence, the above beliefs of Mayor's 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 lies 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 persuades/lies at some stages. For example, consider a game in which the Firm follows the optimal Persuasion strategy in 'normal' times and babbling in 'punishment' times. Punishment periods are triggered whenever a Firm lies about a Bad project and go on long enough to make the Firm indifferent (*modulo* discrete number issues) between lying and telling the truth. In such an equilibrium the Mayor's get zero surplus and the Firm get the same surplus as from Truth Telling. In general, the set of persuasion equilibria that can be supported yield the average payoffs on the line DT. Repeated games make it possible to support some degree of persuasion in equilibrium, but this is Pareto dominated by Truth Telling. It follows then that without a mechanism for generating commitment or for improving on the set of equilibria achievable by repeated games it is best to forego persuasion.

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 project (conditional on being Accepted) is revealed to be either Good or Bad. After this, the Mayor goes to the table, lifts the cup, observes the coin and records whether it was heads or tails.[4] All Mayors now observe the history of reports, project qualities, and coin flips from previous terms. The Firm goes through the same process with each Mayor.

For sufficiently patient Firms there is now a maximally persuasive equilibrium. In this equilibrium 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 all believe the Firm is babbling. The off path threat of the Babbling equilibrium is enough to ensure the Firm never wants to deviate.

This simple transfer free mechanism[5] achieves the maximally persuasive equilibrium by allowing the Mayor to verify whether the Firm is keeping to the prescribed probability of lying. Without such a mechanism the beliefs of subsequent Mayors have to generate on path punishments that make the Firm indifferent between sending an Accept or Reject when a project is bad, otherwise the Firm won't randomize in the prescribed fashion. Now, by introducing the coin and cup the Firm never gets punished on path because the Mayors know whether the Firm was being honest about its level of dishonesty or not. In some sense the coin and cup mechanism bears resemblance to a sunspot, it allows players to know where people are in the game tree so there is no need for inappropriate (on path) punish. However, in this case it is a staggered sunspot, a standard sunspot does not do the job as if both players saw the coin when flipped the game would be changed in only the most trivial of senses.

[4]For those who might worry, the table was in a room with a time triggered lock.

[5]Obviously, in the real world one would probably use an electronic randomization device rather than something so crude as a coin and cup. This would allow further a greater range of policies.

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. His payoffs from action a_t depend on an unknown state of the world, $\theta_t \in \Theta = \left\{\theta^1, \theta^2, \ldots, \theta^N\right\}$, where we denote the cardinality of Θ by N. His payoffs are given by the utility function $u_R\left(a_t, \theta\right)$. Each θ_t is drawn independently, from a prior distribution represented by the vector $\mu_0 \in \Delta^N$.[6] In each period, R_t is ex ante uninformed about θ_t. Each R_t lives only in period $t$and thus in each period t, plays her myopically optimal strategy.

At the beginning of each period, an infinitely-lived Sender S privately observes the realization, θ_t. Before R_t takes an action, S can send a message m_t from some set, M. With some notational abuse, we occasionally use M to refer to the cardinality of the message space. Within a period, the Sender only cares about the action taken by agent R_t and has stage utility $u_S\left(a_t\right)$.

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

1. θ_t is drawn from distribution μ_0. S privately observes this realization.

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, the state θ_t is observed by all players.

So the Sender first observes θ_t and chooses a message in accordance with his optimal strategy. After receiving message m_t, the Receiver R_t forms her posterior belief μ_t using Bayes' rule given the strategy of the Sender; she then chooses his action $a_t\left(\mu_t\right)$ to maximize her expected utility $\mathbb{E}\left[u_R\left(a_t, \theta_t\right) \mid m_t\right]$ and then dies. Hence, we can write the Sender's equilibrium period-t stage payoff, as a function of the Receiver's posterior:

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

As in Kamenica & Gentzkow (2011), we focus on Sender-preferred equilibria. That is, whenever R_t's posterior belief leaves her indifferent between two actions, we assume she chooses the one S prefers. This ensures that $v\left(\mu_t\right)$ is a lower semi-continuous function.

[6] Following standard notation, we use ΔX to denote the simplex over set X, and Δ^N for the N-dimensional unit simplex.

We refer to this stage game by Γ_t. As the stage game is a standard cheap talk game, there always exists a babbling equilibrium. In general, this is not the unique equilibrium of the stage game. Informative equilibria can be sustained so long as the equilibrium messages are all equally profitable for S.[7,8] This can happen, for instance, when $|\Theta| = 2$, and $v(\mu)$ is non-monotonic, with $v(\mu') = v(\mu'')$ for some $\mu' < \mu_0 < \mu''$.

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.

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 meesage space M, and a garbling $\hat{s}$ of θ_t). The key distinguishing feasture 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.

Of course, S can do at least as well using information design as she can in any equilibrium of the static cheap talk game. Define $\hat{v}(\mu)$ as the smallest concave function that is everywhere weakly greater $v(\mu)$. 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

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

[7] Chakraborty & Harbaugh (2010) study the class of equilibria generally attainable in this stage game, under some natural assumptions on payoffs.

[8] Whenever the stage game has multiple equilibria, it will typically be possible to enforce a wider set of outcomes than repeated play of stage equilibria, even for finitely repeated interactions. In these cases, it is nevertheless easier to sustain additional equilibria when the game is infinitely repeated. Moreover, in several highly studies classes of problem babbling is the only equilibrium of the stage game. In these cases, we need infinite repetition to sustain equilibria.

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 history $\phi_t = (m_t, a_t, \theta_t)_{\tau=0}^{t-1}$, the Sender sends a message to a new Receiver R_{τ}. The Receiver observes ϕ_t and the message m_t and then chooses an action a_t before dying. 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 $s_t : \Phi_t \times \Theta \to \Delta M$ 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 weak Perfect Bayesian equilibria of the above game. An equilibrium specifies *(i)* a strategy s for the Sender; *(ii)* strategies $\rho = (\rho_t)_{t=0}^{\infty}$ for each R_t and *(iii)* 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
$$\mathbb{E}\left[\sum_{\tau=t}^{\infty} \delta^{\tau} u_S (a_t) \mid \phi_t, \theta_t; \rho\right].$$

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

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

3.3 A Direct Equilibrium

At any history h_t, a message m sent under the Sender's behavioural strategy $s_t (h_t, \theta_t)$ induces a posterior belief $\mu_t (h_t, m)$ for the Receiver, on which he bases her optimal action. As we have already noted, Sender's payoffs from such a message can be expressed as a reduced-form function of R_t's belief, $v(\mu_t)$. More broadly, $s_t(h_t, \theta_t)$ induces a lottery $\lambda \in \Delta(\Delta\Theta)$ over R_t's posterior beliefs. Of course, since R_t is a rational Bayesian, these induced posteriors integrate back to the prior. A key insight of Kamenica & Gentzkow (2011) is that this restriction is the only constraint on the lotteries λ that can represent some strategy. For any prior μ_0, we refer to the (convex) subset of such feasible lotteries by $\Lambda(\mu_0) \subset \Delta(\Delta\Theta)$. In other words, if at some history h_t, λ induces a lottery of M posteriors $(\mu_{t,1}, \mu_{t,2}, \ldots, \mu_{t,M})$ with $\lambda_j = \Pr(\mu = \mu_{1,j})$, $j = 1, 2, \ldots, M$ that satisfies

$$\mu_0 = \sum_{j=1}^{M} \lambda_j \mu_{t,j} \tag{2}$$

then there exists a strategy $s_t(h_t, \theta_t)$ that could have generated this distribution of posteriors from S's strategy at history h_t.

With these observations in hand, 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 the following adaptation of game Γ_t: $\hat{\Gamma}_t$ 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 $h_t = (\tilde{\mu}_t, a_t, \theta_t)_{\tau=0}^{t-1}$, the set of all period-$t$histories H_t, and (behavioural) strategies functions of the form $\sigma = (\sigma_\tau(h_\tau, \theta_\tau))_{\tau=0}^{\infty}$, where each $\sigma_t : H_t \times \Theta \to \Delta M$, 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

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

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 (3) at all histories of the game tree, given $\mu_\tau(h_\tau, \tilde{\mu}_\tau)$, $\tau \geq t$. Moreover, a *direct equilibrium* requires that *(i)* R_t's beliefs conform to the recommendation made by S, for any $\tilde{\mu}_t$ on the equilibrium path, *(ii)* at any history, S's mixed strategy over messages can be 'averaged back' 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. The intuition for the Lemma is as follows: For any equilibrium in which S uses message $\tilde{m}$ to induce R_t to take an action $\tilde{a}$, it must be the case that R_t's posterior belief $\tilde{\mu}$ made $\tilde{a}$ optimal for him. 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. Of course, under this new messaging strategy, R_t's posterior belief always satisfies the recommendation (that is, condition 2 holds at each history). Moreover these recommendations must be optimal for S, since the underlying equilibrium messages $\tilde{m}$ were optimal in the original game (condition 1 holds at each history). Off the equilibrium path, R_t's beliefs are not constrained by Bayes' rule and thus we can support the equilibrium with beliefs $\mu_t\left(h_t, \tilde{\mu'}_t\right) = \underline{\mu}$, where $\underline{\mu} \in \arg\min v(\mu)$.[9] Finally, as we described above any equilibrium strategy must induce posteriors that satisfy the Law of Total Probability, (2). Since R_t's beliefs conform to $\tilde{\mu}_t$, this same condition carries over to equilibrium messages in game $\hat{\Gamma}^\infty$. Condition 3 simply restates that this must be the case, in terms of requiring that μ_0 live in the convex hull of the support of S's strategy at each history.

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 her optimal discounted average payoffs under

[9]If no minimum exists, we can simply use a function $\mu_t(h_t)$ where $v\left(\underline{\mu}_t\right) = \inf v(\mu) + \epsilon_t$, for some ϵ_t chosen sufficiently small that the current payoff is lower than the worst stage payoff among on-path beliefs (which is well-defined: See Lemma).

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

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.* [10]

While there exist nonconvex 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 $\hat{v}(\mu_{x,y})$, $\forall \mu_{x,y} = \alpha\mu_x + (1-\alpha)\mu_y$, $\alpha \in [0,1]$ (see Figure 4). 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[11]) 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.

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 $\theta \in P$.

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

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

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.

4.1 Repeated Persuasion

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.

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

Proposition 1 generalizes Hermalin (2007), Proposition 1. 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[12]. 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:

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

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

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 $\sigma = (\sigma_1, \sigma_2(h_2), \dots)$ for S that maximizes his present discounted utility, subject to: *(i)* each choice of $\mu_t \in supp(\sigma_t(h_t,\theta_t))$ 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 to be skeptical after such reports.[13] In problem (4), when we choose a strategy σ', we are also able to vary R_t's beliefs following any message sent, so long as they conform to equilibrium restrictions. In addition, the choice of strategy must be optimal for S, given the Receivers' beliefs.

To help understand more properties of the solution to (4), we first present a useful Lemma:

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

[13] 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.

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.[14]

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 equilirium, 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-preferrd one at some θ_t, followed by reversion to the original strategy thereafter. Since the worst cheap talk equilibrium yields statewise

[14]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.

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), 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.

5 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 her 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 occuring 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, $\omega_t \in \Omega$, to the repeated game in Section 3. To ease exposition, we suppose that ω_t is i.i.d. over time $\omega_t \sim U[0,1]$. The CnC mechanism is a combination of sequence of random variables ω_t and repeated play of the following adapted stage game:

1a. θ_t, ω_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, ω_t and a noisy signal of the state θ_t are both 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 ω_t before R_t; *(ii)* Sender cannot manipulate the realization of ω_t; *(iii)* R_t is able to observe ω_t. In this way, one can think of ω_t as a staggered sunspot.

As the next Proposition shows, the CnC mechanism always admits equiliria 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, ω_t), where ω_t does not involve any new payoff considerations for Sender or Receiver. On this expanded state space, we can use ω_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 $\omega_t \in [a, a+\lambda_1]$, $a+\lambda_1 \leq 1$, and μ_2 to be reported if $\omega_t \in [b, b+\lambda_1], b+\lambda_1 \leq 1$, for disjoint intervals $[a, a+\lambda_1]$, $[b, b+\lambda_1]$.[15] Importantly, this means that each round Sender can credibly mimic the optimal experiment, knowing that if he reports according to her '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 babbbling equilibrium.

Importantly, for the CnC mechanism to improve equilibrium outcomes, Sender must know the realization of ω_t *before* reporting (to know which message to send without punishment), while Receiver must only observe ω_t *after* taking action a_t. If Receiver learned ω_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 ω_t to Receivers is a crucial element of the design of CnC mechanisms.

Such staggered sunspots ω_t are implementable and don't require any specialist knowledge of the decision problem or the distribution of θ.[16] 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.[17] The proportion of 'Heads' across the cups can then play the role of ω_t. As a more realistic example, Blockchain technologies (such as that underpinning Bitcoin) support decentralized recording and updating of information

[15] 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.

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

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

among peers using cryptographic methods.[18] 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 prespecified times.[19] Programmed with a random number generator to update the Blockchain with new messages ω_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.

5.1 Financial Advice & Disclsoure Rules

Investment banks and financial brokerages provide several services to clients, helping them to access financial markets: they act as direct trading counterparties, execute trading orders on clients' behalfs and/or provide asset management services on some allocation of funds. Often, these firms are paid via commissions on trades, and ongoing fees on portfolios managed. Moreover, many larger firms employ financial analysts, who conduct research on the markets and provide explicit trading recommendations to clients.

In 2002, the National Association of Security Dealers, a financial industry self-regulating body, imposed rules that required banks to disclose the aggregate distribution of 'Buy', 'Hold' and 'Sell' recommendations their analysts were making across assets.[20] Barber et al. [2006]report that prior to the introduction of these rules, average analyst recommendations were heavily skewed towards Buy recommendations: 60% Buy recommendations; 35% Hold; and only 5% Sell. However, on introduction of these new rules Buy recommendations dropped almost immediately to 51% and buy the following year the ratios were 42% Buy, 41% Hold and 17% Sell. Nor was this merely a drop from an anomylous high, since 1996 Buy recommendations had been greater than 60% and Sell recommendations lower than 5%. Moreover, Kadan et al. (2009) show that prices were more responsive to analyst reports after the introduction of the rules.

Using an adaptation of our model, we can study how the introduction of these disclsoure 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,

[18] 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.

[19] http://www.economist.com/news/science-and-technology/21699099-blockchain-technology-could-improve-reliability-medical-trials-better

[20] For further details, see Kadan et al. (2009).

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.[21] 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.

5.2 Online Platforms & Coarse Reviews

Several online platforms, such as Amazon, eBay, Booking.com, and AirBnB allow consumers to search through listings of independent sellers of goods and services, such as hotels and restaurants. On these platforms sellers can often provide descriptions and photographs advertising their products, to help encourage sales. However, not all experiences for buyers on these sites turn out to be good *ex post*, with some products failing to meet the standards promised on the platform. An obvious way to discipline sellers on these websites is to create reputational concerns by allowing customers to leave feedback about sellers online. However, while platforms make feedback possible, they often aggregate individual feedback into coarse signals of the seller's trustworthiness.

Again, interventions by online platforms to coarsen feedback given by buyers could be interpreted in the light of allowing sellers to persuade effectively. Allowing detailed feedback places sellers in a pure repeated game setting, in which attempts at persuasion would necessarily have to be met with costly punishments - periods in which no potential buyer

[21] Details available on request.

will believe their advertisements, and subsequently shun their prospects. Indeed, in simple binary action model, we have seen that sellers could do no better than in the truth-telling equilibrium. However, if the online platform instead provides information on *only* the aggregate proportion of buyers who were dissatisfied with their purchase (i.e. those who had been falsely persuaded) in an interval of time, this statistic again allows for equilibria in which sellers can achieve payoffs as if they were able to commit to their advetising strategies.

We note that the results of our model apply well when buyers do not have common preferences over a seller's good - i.e. the 'state' variable is a horizontal match between buyer and product, as it might be on AirBnB where location, amenities, quietness, etc. matter. Alternatively, our model applies well if the seller sells multiple different types of good online, as they might on eBay or Amazon. In these settings, others' preferences can play the role of payoff-irrelevant signal, on which aggregate statistics can be used to discipline sellers.

5.3 Financial Roadshows

When investment banks arrange the sale of new debt and equity securities on behalf of an issuing company, a common practice is to arrange tours of presentations ('roadshows') around the world which act as an opportunity to advertise the new assets separately to different members of the investment community. After these roadshows, these assets are sold to the market.

At the time of any individual presentation, an individual investor must both process the information provided about these assets and consider the credibility of the source, comparing the investment bank's previous sales pitches to the subsquent peformance of those assets. However, investors can also condition their beliefs about the credibility of the bank's statements based not only on the individual reports it received, but also indirectly on the reports that other inestors received in the past, through the bidding interest at Initial Public Offerings. In work available on request, we show that the ability to condition beliefs about the bank's future credibility on the reports it sent to all investors (about the same asset) can again be used as the basis to secure optimal persuasion for the investment bank in a repeated cheap talk setting without the need for costly on-path punishments, so long as there are sufficiently many to whom it can sell.

6 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 ways of changing the game so that we can achieve the optimal outcome for the Sender.

These results have important implications for the real world and the literature on persuasion. Perhaps the biggest implication is that it is quite feasible that we have very bad equilibria in the world where people with information are using dishonest persuasion in a fashion that benefits them little but might be very harmful to those relying on their information. If such cases can be identified it may be possible through explicit recognition of the problem and the incentives to move to a better, more honest, equilibrium. However, further to this, when stuck in such a bad equilibrium there are easily created cryptographic mechanisms that allow us to strictly improve the outcomes for both Senders and Receivers. Given this strictly improves the payoffs of both Senders and Receiver they should have a strong incentive to create and coordinate on the use of such a device. Of course, there are also settings where greater ability to commit arises more naturally from repetition. In these cases we should see a greater ability of Senders to persuade and further we should see Senders, and perhaps Receivers, trying to create such settings in the first place.

These results and insights demonstrate a clear tension between honest and dishonest persuasion. On one hand, we show that dishonesty can serve Senders very well when the setting allows for honestly dishonest equilibria. However, on the other, we see that dishonesty can also give the Sender nothing beyond truth telling while worsening the outcomes of Receivers. This suggest that we should take caution when thinking about what kind of advice we give about the extent to which experts should try to persuade through dishonest methods. For if we don't pay attention to the settings and the way such setting can generate commitment we may find such advice doing much more harm than good.

References

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.

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

J. C. Ely. Beeps. Working Paper., 2015.

B. E. Hermalin. Leading for the long term. *Journal of Economic Behavior and Organization*, 62(1):1–19, 2007.

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.

C. Margaria and A. Smolin. Dynamic Communication with Biased Senders. 2015.

E. Perez-Richet. Interim Bayesian persuasion: First steps. *American Economic Review*, 104 (5):469–474, 2014.

L. Rayo and I. Segal. Optimal information disclosure. *Journal of political Economy*, 118(5): 949–987, 2010.

J. Sobel. A Theory of Credibility. *Review of Economic Studies*, 52(4):557–573, 1985.

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

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. To Be Added. □

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.[22]

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 $\Lambda(\mu_0)$.[23] Since v' is upper semicontinuous 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 $\tilde{\Lambda}(\mu_0)$.[24] 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

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

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

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

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 $\lambda \in \Lambda(\mu_0)$. We show that there exists a subsequence $\left\{\mu_0^{n''}\right\}$, such that $\exists \lambda^{n''} \in \Lambda\left(\mu_0^{n''}\right)$ 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 $i \neq j$, $i, j \in \{1, 2, \ldots, N+1\}$.[25] 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 $\left(\left(\tilde{\mu}_0^{n'}, \lambda_0^{n'}\right)\right) \to (\mu_0, \lambda)$ for some $\lambda^\star \in \Lambda(\mu_0)$.[26] 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,$j \in \{1, 2, \ldots, N+1\}$.[27] 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 $\tilde{\epsilon}_2'$, $\delta_1' > 0$. □

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

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

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

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 $m \in M$,

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

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

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

Given any history, we define an equilibrium message $m^\theta \in M'$ to be *uniquely proscribed* at state θ if $supp(\sigma_\tau(h_\tau, \theta)) = \left\{m^\theta\right\}$. 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 $\mu_\tau(\tilde{m}) \in \{\mu_\tau(m)\}_{m \in M'}$ that can be removed from the support such that remaining posteriors still satisfy Bayes' plausibility

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

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 $\mu_\tau(\tilde{m}) \in \{\mu_\tau(m)\}_{m \in M'/\{\tilde{m}\}}$ can be sustained by a feasible signal structure with $N'-1$ distinct messages. Moreover, the message $\tilde{m}$ cannot be uniquely proscribed in any state $\theta \in \Theta$. Otherwise, there would exist some θ^i for which $\mu_\tau^i(m) = 0$, $\forall m \in M'/\{\tilde{m}\}$, while $\mu_0^i > 0$, violating (6). Therefore, $\tilde{m} \in M'/M^\Theta$ and for every state θ in which σ proscribes

$\Pr(m_\tau = \tilde{m} \mid h_\tau, \theta) > 0$, there exists another message m'_θ sent with positive probability in state θ.

Construct a new strategy $\sigma^\star$ which induces the distribution $(\alpha_{m_\tau})_{m_\tau \in M'/\{\tilde{m}\}}$ over the posteriors $\{\mu_\tau(m)\}_{m \in M'}$ at history h_τ, and plays according to σ otherwise (this is feasible, by Proposition 1 of Kamenica & Gentzkow (2011)). For any $m \in M'/\{\tilde{m}\}$, the strategy continues to induce belief $\mu_\tau(m)$ at history h_τ and leaves continuation payoffs unchanged at $V(h_\tau, \theta_\tau)$ thereafter (for any $\theta_\tau \in \Theta_\tau$). Moreover, this continuation payoff is well defined for each m since $\tilde{m}$ was never uniquely 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 compatiblity of any history h_t, for $t > \tau$. But this is a contradiction to N' as the minimum number of messages in any strategy consistent with $\mathbb{E}[V(h_\tau, \theta_\tau)]$. □

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

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

Consider the following problem:

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

s.t.

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

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

$\forall \underline{h}_t$ s.t. $\tilde{\mu}_t = \underline{\mu}_\tau^i$ at all subsequences $\underline{h}_{\tau'}$, $0 \le \tau' \le \tau$, 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.[28]

At $t = 0$ (where $h_0 = \emptyset$), problem (7) is a relaxed version of problem (4): it only retains

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

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 $V_t^\star(\underline{h}_t, \theta_t)$ 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 $\sigma_t(h_t, \theta) \in \Lambda(\mu_0)$ and feasible payoffs from the continuation equilibrium, $V_{t+\tau+1} \in \mathcal{V}$.[29]

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:

$$(1-\delta)\,\mathbb{E}_\theta\left[V_0^\star(\theta_0)\right] = \max_{\lambda \in \Lambda(\mu_0)} \sum \mu_0^i \underline{v}_i(\lambda) \tag{9}$$

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

Since $\underline{v}_i(\lambda)$ is the minimum of finitely many upper semicontinuous functions (from Lemma 2, we need only choose from N distinct posteriors, μ, and v is upper semicontinuous), it is upper semi-continuous. Moreover, the set $\Lambda(\mu_0)$ is clearly compact. Therefore, by the Extreme Value Theorem, the maximum exists. Let the lottery that achieves this optimum be $\lambda^\star \in \Lambda(\mu_0)$, with associated support $\{\mu_1^\star, \mu_2^\star, \ldots, \mu_{N'}^\star\}$, where for convenience we index such that $v(\mu_1^\star) \leq v(\mu_2^\star) \leq \cdots \leq v(\mu_{N'}^\star)$

Finally we show there exists a $\underline{\delta} < 1$ such that for all $\underline{\delta} \leq \delta < 1$, it is possible to construct an equilibrium which attains this upper bound. Specifically, consider the following strategy profile $\sigma^\star$:

- At any history $\underline{h}_t$, S plays $\lambda^\star$ forever, subsequent R_t play their best strategies;[30]

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

[30] Subject to choosing the Sender-optimal action among those she is indifferent between.

- For any history $\left(\underline{h}_{t-1}, \tilde{\mu}_t\right)$, $\tilde{\mu}_t \in supp\left(\lambda^\star\right)$, $\tilde{\mu}_t \neq \underline{\mu}^{i,\star}$, revert to the ***worst possible stage game equilibrium*** for $K_i\left(\tilde{\mu}_t, \delta\right)$ periods, where K_i is chosen to make S indifferent (the final period of this punishment can be implemented with a public randomization device to get around discreteness problems).[31] This punishment is always weakly worse than the solution to (9). If this difference is strict, we can use the worst equilibrium as a punishment for some appropriate $\underline{\delta} < \delta$, for each on-path deviation. If it is weak, then we just use a strategy of playing the worst stage equilibrium forever to implement the payoff - this won't require punishment histories.

- Note that the best 'stage' deviation payoff is uniformly bounded above by $\max_{\mu \in \Delta\Theta} v\left(\mu\right) - v\left(\mu_1^\star\right)$, which is necessarily finite, since $v\left(\mu\right)$ is an upper semi-continuous function on the convex set $\Delta\Theta$. Therefore, we will be able to find a single $\underline{\delta}$ that can sustain punishments appropriately (i.e. we won't need to take δ all the way to 1 to ensure no temptation to cheat at all histories).

- After period K of a punishment phase, return to playing according to $\sigma^\star$ (as if you are at history h_0 - i.e. play $\lambda^\star$ until a deviation, and punish accordingly).

- For any history that is $K' < K$ periods into the punishment, maintain play of the worst static equilibrium in all histories (regardless of 'on-path' deviations within this phase)

- For any off-path deviation $\tilde{\mu}_t$ at any point in the game $\tilde{\mu}_t \notin supp\left(\sigma_t^\star\left(h_t\right)\right)$, R_t believes $\mu\left(\tilde{\mu}_t\right) = \arg\min\left\{v\left(\mu\right) : \mu \in \sigma_t^\star\left(h_t\right)\right\}$ and we revert to play of the worst static equilibrium thereafter. This clearly makes off-path deviations never profitable at any history.

$\square$

Proof of Theorem 2

Proof. (If) Clearly, the optimal discounted average payoff achievable via information design on each Receiver R_t weakly excceds 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, $\hat{v}\left(\mu_0\right)$, can be implemented by a bijection $\hat{s}_P$ between some partition P of Θ to $M := \left\{m_1, m_2, \ldots, m_{N'}\right\}$, where $N' \leq N$. Thus, for each $\theta^i \in \Theta$, $\hat{s}\left(\theta^i\right) = m\left(\theta^i\right)$, for some unique $m \in M$. Moreover, we can define an inverse function

[31] We choose to enforce punishments with public randomization only for the sake of ease of exposition of the strategy. We can esablish a similar result without public randomization, for the limit as $\delta \to 1$.

$m^{-1}(m_j) := \{\theta : m(\theta) = m_j\} \subset \Theta$, 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 Recevier'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

$$\underline{v}_i(\lambda^{\star\star}) \leq \sum_j \frac{\lambda_j^{\star\star}\mu_j^{i,\star\star}}{\mu_0^i} v\left(\mu_j^{\star\star}\right) \tag{10}$$

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

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

for some beliefs $\hat{\hat{\mu}}_l$ in the support of ΔP_j.[32] 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 $\Gamma_{P_j}(\mu(P_j))$.

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

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

and therefore

$$\sum \mu_0^i \underline{v}_i(\lambda^{\star\star}) < \sum_{j \in \{1,2,\ldots,N'\}} \lambda_j^{\star\star} v\left(\mu_j^{\star\star}\right) = \hat{v}(\mu_0) \tag{11}$$

A similar argument establishes that the payoff from *any* experiment inducing arbitrary

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

lottery λ' attains a weakly higher payoff than $\sum \mu_0^i \underline{v}_i(\lambda')$ evaluated at λ'. Therefore,

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

where $\lambda^\star$ solves (9). The inequality is generically strict since either: *(i)* $\lambda^\star$ is not an honest strategy (in which case (11) holds generically), or *(ii)* $\lambda^\star$ is honest, in which case $\sum_{j\in\{1,2,\ldots,N'\}} \lambda_j^\star v\left(\mu_j^\star\right) < \hat{v}(\mu_0)$ by our assumption that $\hat{v}(\mu_0)$ cannot be implemented by an honest experiment. $\square$