

Goodwill in communication [☆]

Aditya Kuvalekar ^{a,*}, Elliot Lipnowski ^b, João Ramos ^{c,d}

^a University of Essex, United Kingdom of Great Britain and Northern Ireland

^b Columbia University, United States of America

^c University of Southern California, United States of America

^d Queen Mary University of London, United Kingdom of Great Britain and Northern Ireland

Received 1 May 2021; final version received 5 January 2022; accepted 18 April 2022

Available online 2 May 2022

Abstract

An expert advises a decision maker over time. With both the quality of advice and the extent to which it is followed remaining private, the players have limited information with which to discipline each other. Even so, communication in and of itself facilitates cooperation, the relationship evolving based on the expert's advice. We show a formal equivalence between our setting and one of cheap talk with capped money burning, enabling an exact characterization (at fixed discounting) of the expert's attainable payoffs. While an ongoing relationship often helps, our characterization implies that, absent feedback, relational incentives can never restore commitment.

© 2022 Elsevier Inc. All rights reserved.

JEL classification: C73; D82; D83

Keywords: Cheap talk; Communication; Persuasion; Repeated game

[☆] Lipnowski acknowledges support from the National Science Foundation (grant SES-1730168), and Kuvalekar acknowledges support from the Ministerio Economía y competitividad (grant PGC2018-09159-B-I00) and Comunidad de Madrid, MadEco-CM (grant S2015/HUM-3444) that he received while he was at Universidad Carlos III de Madrid. We would like to thank Dan Barron, Ben Brooks, Odilon Câmara, Joyee Deb, Inga Deimen, Wouter Dessein, Andrew McClellan, Mike Powell, Daniel Rappoport, and especially Doron Ravid.

* Corresponding author.

E-mail addresses: a.kuvalekar@essex.ac.uk (A. Kuvalekar), e.lipnowski@columbia.edu (E. Lipnowski), j.ramos@qmul.ac.uk (J. Ramos).

In many economic relationships, those with authority over a decision differ from those with the relevant expertise to make that decision, and their interests may differ. A sensible means to informed decision-making is for the parties to communicate, as in the model of Crawford and Sobel (1982) (hereafter CS). However, conflicting interests hinder communication: when information transmission is strategic and unverifiable (i.e., “cheap talk”) the expert cannot commit to convey her information truthfully and the decision maker cannot commit to how he will use the advice. For instance, when a policy analyst with domain-specific expertise advises a policy-maker, the latter has little means to assess the quality of the advice before acting; and when a financial analyst advises a client on investment decisions, the client retains authority over how to actually invest the funds.

A hallmark feature of advisory relationships is that they are often dynamic, enabling the players to provide incentives via future rewards to discipline today’s behavior. To provide such incentives, the parties might rely on *feedback* concerning others’ past behavior. For example, to punish advice not given in good faith, the policymaker might use some ex-post feedback about the information on which the analyst based her advice. Similarly, to know whether her advice was followed, the analyst might use some feedback about the policymaker’s choices following her advice.

The premise that motivates this study is that the above feedback may not always be readily available. Feedback about contemporaneous information and play could be too noisy, delayed, or complex to be meaningfully relied upon for sustaining cooperation. Stepping further outside the model, the exact nature of what an expert knows or what options a decision maker faces may be prohibitively difficult for others to assess. Nevertheless, an ongoing advice-based relationship is still an inherently dynamic relationship, because *the advice itself* serves as a shared history between the players. With this perspective, as a theoretical benchmark, we study a repeated game between a sender and a receiver in which the only source of feedback to the players is the sender’s advice.

We observe in this paper that communication in and of itself facilitates cooperation. The key observation is that some advice (when expected to be followed) is inherently more tempting for some types of experts to give. Accordingly, by varying the terms of the relationship based on what the expert suggests, different forms of advice can be *priced*. We formalize this intuition and characterize the limits of this channel.

Formally, we study an infinite-horizon discounted repeated game between an expert and a decision maker. Each period, an independent payoff state is drawn and privately observed by the expert. The expert sends a message to the decision maker who, upon hearing the message, privately chooses an action. As the decision maker never observes the (current or past) state, and the expert never observes the decision maker’s (current or past) action, neither is used to vary the future terms of the relationship. Even so, we show that conditioning future play on today’s communication can aid cooperation. To trace the limits of such cooperation, we ask which expert payoffs can arise in equilibrium.

In order to characterize equilibrium outcomes, we consider two auxiliary static games. First, we consider our stage game (with identical action space for the decision maker, state distribution, and payoffs for both players), further endowing the expert with the ability to observably *burn money* along with any message, up to an exogenously specified cap $M \geq 0$. The expert’s highest equilibrium payoff in the auxiliary game is weakly higher than in our true stage game and, as Austen-Smith and Banks (2000) demonstrated for the leading example of CS, can be strictly

higher. Dually, we consider our stage game with the expert further allowed to observably *collect bonuses* along with her sent message, up to a cap.¹

Our main result is that the set of expert equilibrium payoffs in our repeated game exactly coincides with those attainable in these auxiliary static games (with the cap being the largest one compatible with a dynamic enforcement constraint). This theorem, which is nearly an immediate consequence of the standard recursive toolbox (e.g. Abreu et al., 1990, hereafter APS),² still provides new insights for long-run advisory relationships. First, conceptually, we formalize a sense in which the future relationship can be used to price an expert's various recommendations. Next, we make progress toward the question of which types of advisory relationships are served by relational incentives; in particular, the theorem readily implies that repetition can never help an expert whose preferences are independent of the state. Further, our main result implies that a very patient expert's best equilibrium payoff is her best payoff from static cheap talk with (uncapped) money burning; repetition can therefore never restore her "Bayesian persuasion" value. Finally, we apply our characterization to a repeated game of project funding, solving for the expert's equilibrium payoff set at every discount factor by analyzing a textbook, static mechanism problem.

Related literature

Since CS and Green and Stokey (2007), sender-receiver games featuring "cheap talk" communication have been a canonical framework for studying information transmission between an informed expert and an uninformed decision maker. More broadly, a large literature studies communication games with other protocols, including mediated communication (e.g. Forges, 1986; Salamanca, 2021), back-and-forth communication (e.g. Forges, 1990; Krishna and Morgan, 2004; Aumann and Hart, 2003), communication with evidence (e.g. Glazer and Rubinstein, 2006; Hart et al., 2017; Rappoport, 2021), flexible communication with commitment (e.g. Aumann and Maschler, 1966; Kamenica and Gentzkow, 2011), and more. Our work, of course, belongs to this broader literature.

More specifically, we join a very active literature on strategic communication in long-run relationships, asking to what extent dynamics enable effective communication.³ The literature has studied the possibilities for effective advisory relationships when the advice of the expert can be assessed after the fact (Best and Quigley, 2020; Mathevet et al., 2019); when the decision maker's behavior can be observed and disciplined (Alonso and Matouschek, 2008; Renault et al., 2013; Margaria and Smolin, 2018; Kolotilin and Li, 2021), and when decisions are interspersed with advice about a persistent state (Golosov et al., 2014). We study a model in which the relationship is ongoing but, even with patient players, the core tension of a cheap talk interaction remains: The expert faces no dynamic consequences from being caught lying—for she is never caught—and the decision maker faces no dynamic consequences from acting to his own myopic benefit.

¹ Formally, we consider attainable payoffs from cheap talk with capped bonuses, as we range over all positive caps $\tilde{M} \leq M$.

² The only technical hurdles to the analysis concern properties of the auxiliary static communication game. Once these properties are proven, the equivalence result follows immediately from APS's arguments. We refer the reader to the discussion in section 2.2 for more details.

³ Less directly related is work that studies how communication can help sustain cooperation in long-run relationships facing no asymmetric information about payoff-relevant parameters (e.g. Compte, 1998; Kandori and Matsushima, 1998).

Our central result is that a repeated cheap talk game is, in terms of expert equilibrium payoffs, formally equivalent to a conceptually simpler class of games: static cheap talk games with money burning or bonuses. Static cheap talk games with the sender burning money have received some attention in the context of the one-dimensional CS model, starting with the important observation (Austen-Smith and Banks, 2000) that the possibility of burning money can strictly expand the scope for equilibrium communication; relatedly, Krishna and Morgan (2008) discuss how transfers should optimally be used to elicit information in the same setting. Capped money burning makes an appearance in this literature as well, with Kartik (2007) showing that equilibrium communication with severely capped money burning is very close to equilibrium communication with cheap talk alone. More recently, Karamychev and Visser (2017) study the sender-optimal equilibrium of the leading example of CS with money burning; their work is (given our results) an important input to the computation of the expert's equilibrium payoff set in the repeated version of this canonical example. At the end of Section 2, we explain how to use their solution to build such an optimal equilibrium in the associated repeated game.

A key ingredient of our analysis is the standard recursive toolbox for repeated games, à la APS.⁴ As we recursively study the payoffs of one player when our other player cannot be provided dynamic incentives, a central pair of references is Fudenberg et al. (1990) and Fudenberg and Levine (1994), which study repeated games with a mix of long- and short-run players. While their chief focus is on the limit as long-run players become patient, our analysis is still indebted to their crucial observation that the long-run players' payoffs are amenable to recursive analysis.

Our main theorem is reminiscent of results in the literature on relational contracts with transfers, (e.g. Levin, 2003). In that literature, transferable utility typically makes it tractable to quantify the extent to which the future value of the relationship enables cooperation—through a so-called dynamic enforcement constraint (a.k.a. self-enforcement constraint). Kolotilin and Li (2021) use the same to shed light on relational communication with transfers, when the decision maker's play can be perfectly monitored. The presence of transfers and the observability of the decision maker's actions enable the expert to reward the decision maker for compliant decisions and also allows her to credibly signal her private information. While the players in our repeated game do not have access to transfers, money burning, or any other kind of monetary incentives, we show how the expert's payoffs from our repeated game are tightly linked to her payoffs from auxiliary static games with monetary incentives.

1. Model

We consider an infinite-horizon game played between an expert (the sender, S, she), and a decision maker (the receiver, R, he). Time is discrete and indexed by $t \in \{0, 1, 2, \dots\}$. Players S and R play the same stage game in every period, and each player $i \in \{S, R\}$ discounts the future at the rate $\delta_i \in [0, 1]$; let $\delta := \delta_S$.

Each period t , players first observe a sunspot (i.e. a public randomizing device) with uniformly distributed outcome $\omega_t \in [0, 1]$. Then, a state $\theta_t \in \Theta$ is realized, following full-support prior $\mu_0 \in \Delta\Theta$ —and is observed privately by S. Following this, S sends a public message $z_t \in Z$ to R. Having received the message, R privately chooses an action $a_t \in A$, and each player $i \in \{S, R\}$ accrues flow payoff $u_i(a_t, \theta_t)$. Crucially, at the end of the period, the only public information

⁴ Our work is therefore also related to the vast literature in repeated games that applies said recursive toolbox, some of which is surveyed in Mailath and Samuelson (2006). Many such studies also draw an analogy between monetary incentives and variation in continuation payoffs, Fudenberg et al. (1994) being a prominent example.

available is the message that was sent (together with the public sunspot).⁵ The family $\{\omega_t, \theta_t\}_{t=0}^\infty$ of random variables are all independently distributed. Each player $i \in \{S, R\}$ seeks to maximize $\mathbb{E} \sum_t (1 - \delta_i) \delta_i^t u_i(a_t, \theta_t)$, which we spell out more formally below.

Finally, we make some technical assumptions. The spaces Θ , A , and Z are assumed compact metrizable, and the objectives $u_S, u_R : A \times \Theta \rightarrow \mathbb{R}$ continuous.⁶ Moreover, to rule out exogenous (i.e., non-incentive-based) frictions to communication, we assume Z is uncountable.⁷

1.1. Histories, strategies, and equilibrium

For any $t \in \mathbb{Z}_+$, a time- t history is an element of $\mathcal{H}_t := ([0, 1] \times \Theta \times Z \times A)^t \times [0, 1]$. Thus, a history records each past period's sunspot, state, message, and action, in addition to the current period's sunspot. Let $\mathcal{H} := \bigcup_{t=0}^\infty \mathcal{H}_t$ denote the set of all histories. For each history $h \in \mathcal{H}_t \subseteq \mathcal{H}$, the induced private histories for S and R are given by its projections onto $\mathcal{H}_t^S := ([0, 1] \times \Theta \times Z)^t \times [0, 1]$ and $\mathcal{H}_t^R := ([0, 1] \times Z \times A)^t \times [0, 1]$, respectively. Defining $\mathcal{H}^S := \bigcup_{t=0}^\infty \mathcal{H}_t^S$ and $\mathcal{H}^R := \bigcup_{t=0}^\infty \mathcal{H}_t^R$, we are now equipped to define the ingredients of an equilibrium. A strategy for S is a measurable map $\sigma : \mathcal{H}^S \times \Theta \rightarrow \Delta Z$; a strategy for R is a measurable map $\rho : \mathcal{H}^R \times Z \rightarrow \Delta A$; and a belief map for R is a measurable map $\beta : \mathcal{H}^R \times Z \rightarrow \Delta \Theta$. The interpretation is that (following history h): S [resp. R] will use contemporaneous message [resp. action] distribution $\sigma(\cdot | h^S, \theta)$ [resp. $\rho(\cdot | h^R, z)$] if the current state [resp. message] is θ [resp. z]; and R will have belief $\beta(\cdot | h^R, z)$ if the current message is z .⁸

For any time $t \in \mathbb{Z}_+$, current history $h \in \mathcal{H}_t$, and current state $\theta = \theta_t \in \Theta$, a strategy profile (σ, ρ) defines a distribution over $(\omega_\tau, \theta_\tau, z_\tau, a_\tau)_{\tau=t}^\infty \in ([0, 1] \times \Theta \times Z \times A)^\infty$.⁹ Letting $\mathbb{E}^{h, \theta, \sigma, \rho}$ denote expectation with respect to this law, player $i \in \{S, R\}$ has continuation payoff given by

$$\Pi_i(\sigma, \rho | h^i, \theta) = (1 - \delta_i) \sum_{\tau=t}^\infty \delta_i^{\tau-t} \mathbb{E}^{h, \theta, \sigma, \rho} \left[u_i(a_\tau, \theta_\tau) \middle| h^i, \theta \right].$$

Definition 1. A (perfect Bayesian) equilibrium is a triple $\langle \sigma, \beta, \rho \rangle$ such that, $\forall h \in \mathcal{H}$:

- (1) S rationality: Every $\theta \in \Theta$ has $\sigma \in \arg\max_{\tilde{\sigma}} \Pi_S(\tilde{\sigma}, \rho | h^S, \theta)$;
- (2) R rationality: Every $z \in Z$ has $\rho \in \arg\max_{\tilde{\rho}} \int_{\Theta} \Pi_R(\sigma, \tilde{\rho} | h^R, \cdot) d\beta(\cdot | h^R, z)$;
- (3) Bayes: Any two $h^R, \tilde{h}^R \in \mathcal{H}^R$ that differ only in the past actions chosen by R have $\beta(h^R, z) = \beta(\tilde{h}^R, z)$; and every Borel $\hat{\Theta} \subseteq \Theta$, $\hat{Z} \subseteq Z$ have

⁵ In particular, we make the extreme assumption that players do not observe their payoffs in real time. This assumption is common in the repeated games literature (e.g. some examples focusing on the communication of private information are Aumann et al., 1995; Renault et al., 2013; Kolotilin and Li, 2021). One interpretation is that realized payoffs are too noisy or delayed to be useful on a realistic timescale.

⁶ In this paper, we endow every compact metrizable space Y with the Borel algebra, and endow the space $\Delta Y = \Delta(Y)$ of Borel probability measures on Y with the weak*-topology.

⁷ In the case that Θ is finite, the S equilibrium payoff set would be the same under the weaker assumption that $|Z| \geq |\Theta|$.

⁸ For a given history $h \in \mathcal{H}$ and player $i \in \{S, R\}$, let h^i denote the projection of h onto $\mathcal{H}^i$. In a mild abuse of notation, any of σ, ρ, β can take an argument from $\mathcal{H}$ rather than $\mathcal{H}^S$ or $\mathcal{H}^R$, with the measurability restriction that it be measurable with respect to the projection.

⁹ Formally, one can recursively define the law for times $t \leq \tau \leq T$ for each T , and then apply Kolmogorov's extension theorem.

$$\int_{\hat{\Theta}} \sigma(\hat{Z}|h^S, \theta) d\mu_0(\theta) = \int_{\Theta} \int_{\hat{Z}} \beta(\hat{\Theta}|h^R, z) d\sigma(z|h^S, \theta) d\mu_0(\theta).$$

In such an equilibrium, the induced **S payoff** is $v = \int_0^1 \int_{\Theta} \Pi_S(\sigma, \rho|\omega, \cdot) d\mu_0 d\omega$.

For the remainder of the paper, our primary object of interest will be the set V_δ of equilibrium S payoffs (as will become apparent later, V_δ depends on δ but not on δ_R , hence the notational choice). We will give an interpretable characterization of V_δ , and then go on to show that this characterization yields new insights for the theory of advisory relationships.

Remark 1. Notice that R's action is completely private. Hence, repeated interaction cannot be used to discipline R, and he will always choose the myopically optimal action. Our model is therefore equivalent to one with a long-run sender and an infinite sequence of short-lived receivers, in which receivers see the entire history of past messages. Interpreting receivers as short-lived may be better-suited to some applications. For example, it may be natural to study a game of repeated mass-communication to a population of receivers. Given the strategic equivalence, we do not take a strong stance, though our language will reflect the one-receiver model in what follows.

Remark 2. If we augmented our game with any sort of feedback—public or private, immediate or delayed, perfect or imperfect—about past states or actions between periods, the strategies in the present model could naturally be viewed as special strategies in the richer game, namely, those in which players choose not to condition on this additional feedback. Following standard reasoning (analogous to repeating a stage-game equilibrium being an equilibrium in a repeated game), our equilibria would remain equilibria in the augmented game. For this reason, our analysis and results admit an alternative interpretation: We characterize the set of sender payoffs that can be attained in a way that is robust to the form and quality of feedback players receive as their relationship progresses.

Remark 3. Our results focus on the set V_δ of attainable S payoffs in our game. Especially since only S can be provided with dynamic incentives, we see this set as the most natural target to assess the scope for dynamic incentives in the repeated communication game. Having said that, attainable R payoffs may be of further interest. Our results speak to this question in two ways. First, natural examples exist for which the set of equilibrium payoff profiles is one-dimensional. For example, with symmetric discounting in the quadratic-loss, constant-bias specification of CS, it is straightforward to deduce that the sender's and the receiver's payoff differ by a constant. Hence, characterizing S equilibrium payoffs in the repeated game serves to characterize the entire set of equilibrium payoff profiles. Second, more generally, characterizing V_δ is a natural first step toward characterizing the entire set of equilibrium payoff profiles. Indeed, standard reasoning à la Spear and Srivastava (1987) delivers a recursive characterization of the upper and lower boundaries of the set of payoff profiles, both of which are functions with domain V_δ .

2. Pricing advice

Our main theorem is an equivalence result between our repeated game and an interpretable class of one-shot games related to our stage game. In this section, we define the latter class of games more formally, and then state and provide intuition for the equivalence result.

2.1. Money burning and bonuses

The key ingredients of our characterization are one-shot games of cheap talk with money burning and bonuses. The parameters of these static games are those of our original stage game, together with a cap $M \geq 0$. S learns the state $\theta \in \Theta$ drawn according to μ_0 , S sends a message $z \in Z$ and chooses an amount of money $y \in \mathbb{R}$ subject to a restriction, and R (having seen y and z) chooses action $a \in A$; the resulting payoffs to S and R are $u_S(a, \theta) - y$ and $u_R(a, \theta)$. Here, y must lie in $[0, M]$ in the case of M -capped money burning and in $[-M, 0]$ in the case of M -capped bonuses.

We should emphasize that S has no technology to burn money or collect bonuses in our repeated model. Rather, these auxiliary games serve as a solution method for our discounted repeated game. We use money burning to investigate whether S can achieve a higher payoff in equilibrium than what she could in the stage game. Analogously, we can use bonuses to look for equilibria in which S suffers worse equilibrium outcomes than he could in the one-shot game alone.¹⁰

Definition 2. Given closed $Y \subseteq \mathbb{R}$, say $v \in \mathbb{R}$ is **attainable with Y -monetary signals** if there exist measurable maps $\sigma : \Theta \rightarrow \Delta(Y \times Z)$, $\rho : Y \times Z \rightarrow \Delta A$, and $\beta : Y \times Z \rightarrow \Delta \Theta$ such that:

- S rationality: $\forall \theta \in \Theta$,

$$\sigma \left(\operatorname{argmax}_{(y,z) \in Y \times Z} \left[\int_A u_S(\cdot, \theta) d\rho(\cdot | y, z) - y \right] \middle| \theta \right) = 1;$$

- R rationality: $\forall y \in Y, z \in Z$,

$$\rho \left(\operatorname{argmax}_{a \in A} \int_{\Theta} u_R(a, \cdot) d\beta(\cdot | y, z) \middle| y, z \right) = 1;$$

- Bayes: $\forall$ Borel $\hat{\Theta} \subseteq \Theta, \hat{Z} \subseteq Z, \hat{Y} \subseteq Y$,

$$\int_{\hat{\Theta}} \sigma(\hat{Y} \times \hat{Z} | \cdot) d\mu_0 = \int_{\Theta} \int_{\hat{Y} \times \hat{Z}} \beta(\hat{\Theta} | \cdot) d\sigma(\cdot | \theta) d\mu_0(\theta);$$

- Value v :

$$\int_{\Theta} \int_{Y \times Z} \left[\int_A u_S(\cdot, \theta) d\rho(\cdot | y, z) - y \right] d\sigma(y, z | \cdot) d\mu_0 = v.$$

In particular, given $M \geq 0$, say v is **attainable with M -capped money burning** [resp. **attainable with M -capped bonuses**] if it is attainable with $[0, M]$ -monetary [resp. with $[-M, 0]$ -monetary] signals. We say v is **attainable with (uncapped) money burning** if it is attainable with $\mathbb{R}_+$ -monetary signals.

¹⁰ See Appendix A.4 for an example in which an equilibrium of a game with capped bonuses is worse for S than any equilibrium of our stage game. Thus, given our equivalence theorem, repetition can also expand the scope for low S payoffs in our class of games.

Finally, say an interval $[v_0, v_1] \subset \mathbb{R}$ is **attainable with cap M** if some $v \geq v_1$ is attainable with M_1 -capped money burning for some $M_1 \in [0, M]$, and some $v \leq v_0$ is attainable with M_0 -capped bonuses for some $M_0 \in [0, M]$.

Three observations are immediate concerning the last part of the above definition. First, raising the cap will obviously preserve attainability of a given interval. Second, there is a sufficiently large cap $\bar{M}$ such that the exact same payoff intervals are attainable for all caps above $\bar{M}$. Indeed, this follows because the payoffs in our game (gross of money burnt or bonuses collected) are bounded, which readily delivers (by S incentives) a bound on the amount of money burned or bonuses forgone in an equilibrium.¹¹ Third, it is always without loss to take $M_1 = M$. Indeed, from any equilibrium with M_1 -capped money burning, an outcome-equivalent equilibrium can be constructed in which any amount above M_1 that S burns is ignored.¹²

2.2. Main characterization

Our main result details how analyzing the one-shot game with capped money burning and bonuses enables a characterization of the S payoff set in our repeated game. One looks for the largest interval which is attainable with cap M , letting the cap be the largest one compatible with a *dynamic enforcement* constraint.

Theorem 1 (Equivalence theorem). *The equilibrium payoff set V_δ is the largest interval $[\underline{v}_\delta, \bar{v}_\delta]$ such that $[\underline{v}_\delta, \bar{v}_\delta]$ is attainable with cap $\frac{\delta}{1-\delta}(\bar{v}_\delta - \underline{v}_\delta)$.*

The reasoning behind the theorem, formally proven in the appendix, is simple. First, because R acts privately, the repeated interaction cannot be used to provide him with dynamic incentives. S's equilibrium payoff set is therefore identical to what it would be if R were myopic. One can then recursively characterize the set of S equilibrium payoffs, with no need to track the value delivered to R.¹³ Next, given public randomization, we need only characterize the best and worst equilibrium payoffs for S. Consider a best equilibrium payoff for S. Each message she sends on path will result in some continuation value that is, by fiat, weakly less than her total equilibrium value. Relabeling this payoff loss as an amount of money that S burns while sending said message (renormalized to account for discounting), all S incentive constraints in the repeated game are identical to those in the stage game augmented with money burning, and S attains the exact same payoff. Finally, the boundedness of the payoff set from below puts a cap on how much money can be burned under this relabeling. Analogous reasoning ties S's lowest equilibrium payoff to one from a one-shot game with bonuses, and the theorem follows. At the end of this section, we return to more explicitly describe an equilibrium construction that generates a best and worst equilibrium for the S in the repeated game.

As mentioned above, the proof of our main result is, conceptually, a direct descendent of classic results of APS. We should note, however, that some new arguments are required along

¹¹ For instance, $\bar{M} := \max_{a, a' \in A, \theta \in \Theta} [u_S(a, \theta) - u_S(a', \theta)]$ would have this feature.

¹² In contrast, it is not without loss to take $M_0 = M$. For example, in the range where $M > \bar{M}$, a strictly worse sender payoff can arise in equilibrium of the game with $\bar{M}$ -capped bonuses than with M -capped bonuses. The reason for the asymmetry is that superfluous money burning that R ignores is never profitable to S, whereas choosing a higher bonus can serve as a profitable deviation for S if R ignores it.

¹³ This observation is identical to that in Fudenberg et al. (1990).

two dimensions. First, one must show that the appropriate set operator—which maps a set of attainable continuation values for an expert from tomorrow onward to a resulting set of attainable continuation values from today onward—takes compact sets to compact sets (which is immediate in the finite-state, finite-action case). Having established said closure property, a recursive characterization à la APS is obtained with no conceptual novelty whatsoever. The upshot is an equivalence between the repeated game and a one-shot game with communication *and contracting*, in which each message entails an amount money that will be automatically burned once said message is sent.

The second new argument one requires, therefore, is that such commitment power does not expand the scope for profitable communication, i.e. that communication with *discretionary* money burning (or bonuses) is payoff-equivalent to communication with contractually burned money.

The key machinery that enables us to complete these two steps is a novel characterization of equilibrium outcomes in static communication games (see Appendix A.1), somewhat analogous to the equivalence between direct mechanisms and delegated sets in the literature on optimal delegation (e.g. Holmström, 1982). In short, we directly characterize which sets of interim outcomes (i.e. triples of burned money, decision maker belief, and mixed action) arise as the set of inducible-on-path interim outcomes for some equilibrium, and further show how the expert's equilibrium payoff can be directly computed from such a set.¹⁴ This characterization of attainable static outcomes, which we believe to be new, may be of broader interest for the theory of strategic information transmission.

Adapting the APS argument in the obvious way, Theorem 1 also yields an “algorithmic” characterization of V_δ . More importantly, though, this simple result yields new qualitative insights on the nature of advisory relationships, and provides a conceptual simplification that allows us to study concrete examples.

Equilibrium construction. To conclude the section, let us briefly describe the S-optimal equilibrium behavior implicit in the constructive proof of Theorem 1. Let v_1 and v_0 denote the highest and lowest equilibrium S payoff, respectively, and let $M := \frac{\delta}{1-\delta}(v_1 - v_0)$. Let us informally describe a two-state automaton equilibrium in which S obtains a payoff v_ξ when play starts in $\xi \in \{0, 1\}$; in particular, starting in state $\xi = 1$ will yield an S-optimal equilibrium. In order to describe the equilibrium, it suffices to describe the current play in each state, along with the rule for (stochastically) switching between the two states.

To construct the stage play in state 1, note (by Theorem 1) that some equilibrium of the static game with M -capped money burning yields an S payoff of v_1 . Moreover, as we show, the equilibrium can be chosen so that every message z admits a unique amount $m(z)$ of money that S chooses to burn whenever sending message z . Let us construct the stage play in state 1 from the equilibrium of the auxiliary game. Let the S strategy send the same message distribution in each state as the S strategy did from the auxiliary game; and let R's beliefs and behavior in response to a message z be exactly what it would have been in the auxiliary game had S sent message z and additionally burned $m(z)$. Finally, the switching rule from state 1 is also constructed from the auxiliary game's equilibrium. Specifically, whenever a message z is sent, the state switches with probability $\frac{m(z)}{M}$, and remains in state 1 with complementary probability. The construction

¹⁴ While our characterization, when specialized to the case of finitely many states and actions, is not identically the one-shot analysis in Aumann and Hart (2003), the resulting characterization of sender payoff vectors is readily derived from theirs in this simpler special case of the model.

of the stage play and switching probabilities in state 0 is analogously built from an equilibrium of the auxiliary static game with capped bonuses.

An important special case of our model is one in which allowing for bonuses cannot hurt S—that is, in which every equilibrium of a static game with capped bonuses is weakly better for S than some equilibrium of the pure cheap talk game. For example, this class of games includes every example (like that of Section 4) in which R has some best response to the prior that gives S her minimum payoff in every state, and includes every example (like the quadratic CS specification) in which S has Blackwell-increasing preferences over how informed R is. In this special case of the model, because no bonuses are used in the auxiliary static game for state 0, it follows that the latter is absorbing in the automaton we construct. Hence, in our constructed S-optimal equilibrium, some (almost-surely finite) time τ exists such that the same stage-game equilibrium is played indefinitely from period τ onward.

As an example, consider the quadratic-uniform specification of CS. In this example, Karamychev and Visser (2017) provide an explicit characterization of an S-optimal equilibrium of the game with uncapped money burning; let v_1 denote the associated S payoff, which is strictly larger than her payoff v_0 from a babbling equilibrium. The amount of money burned in their constructed equilibrium is bounded above, and so lies below $M = \frac{\delta}{1-\delta}(v_1 - v_0)$ for large enough δ . For such δ , an S-optimal equilibrium is simple to describe. Players begin in state 1, wherein play corresponds to that of Karamychev and Visser's (2017) equilibrium period by period, and the hazard rate of switching to state 0 is proportional to the amount of money that would have been burned in the auxiliary static game; then, once in state 0, a permanent babbling equilibrium is played. Finally, in light of Remark 3, this equilibrium is in fact Pareto dominant when discounting is symmetric. In this constructed equilibrium, the quality of information provided stochastically degrades over time as S depletes R's goodwill toward her.

3. Consequences for advisory relationships

In the previous section, we developed a mathematical link between our discounted repeated advisory game and static games. One purpose of that exercise was to formalize how an expert's future goodwill can be used to price today's advice. As an additional benefit, we show how this formal link delivers new, general lessons about the nature of repeated communication games.

For each of the formal results in this section—all easy corollaries of our main theorem—we provide only a sketch of the straightforward proof.

3.1. Patient experts

Our focus has been on how much cooperation can be sustained on the back of S's evolving goodwill at fixed discounting, and on how this varies with the discount factor. However, it is worth noting a simple consequence of the result: an interpretable description of the limit as players become patient (or interact frequently).

Corollary 1 (*No cap for patient experts*). *There exists $\bar{\delta} \in [0, 1)$ such that, whenever $\delta \geq \bar{\delta}$, the highest equilibrium S payoff $\bar{v}_\delta$ is equal to the highest S payoff attainable with uncapped money burning.*

It is immediate from our theorem that the highest S payoff $\bar{v}_1$ from uncapped money burning is an upper bound on $\bar{v}_\delta$ for every $\delta \in (0, 1)$, so we need only argue that $\bar{v}_1$ is an equilibrium value

for sufficiently high discount factors. There is nothing to show if $\bar{v}_1 = \bar{v}_0$, since repeating an equilibrium from the stage game is always an equilibrium of the repeated game; so assume now that $\bar{v}_1 > \bar{v}_0$. Consider an S-preferred equilibrium with money burning (which our analysis in the appendix shows to exist), and recall that S incentives imply the same is in fact an equilibrium with $\bar{M}$ -capped money burning for some finite $\bar{M}$. The corollary follows from noting that $\frac{\delta}{1-\delta}(\bar{v}_1 - \bar{v}_0) \geq \bar{M}$ when δ is sufficiently close to 1.

Observe, Corollary 1 is not a folk theorem. The corollary states that S's equilibrium payoff set exactly attains its asymptotic limit for some sufficiently large $\delta < 1$, which folk theorems with imperfect monitoring (e.g. Fudenberg et al., 1994) do not generally guarantee. Moreover, the corollary delivers a specific interpretable form of this limit set, which is *not* generally equal to the set of feasible and individually rational payoffs. Indeed, Corollary 3 below implies that, in most interesting cases, these two sets differ.

We note here an important consequence for the applied literature on strategic communication. Models of static cheap talk have been used to study organizational decision making (e.g. Dessein, 2002; Che et al., 2013), legislative bargaining (e.g. Gilligan and Krehbiel, 1989), electoral competition (e.g. Kartik and Van Weelden, 2018), and more. In parallel, the theoretical literature has studied the ways in which the ability to burn money expands the scope for credible communication. Given the above corollary, to the extent that we think relationships between employees, legislators, or politicians are typically ongoing with frequent interactions, the model of cheap talk with money burning might be a more appropriate static model to employ—with money burning serving as a reduced form for degradation of the future relationship. Further, given our conservative model of goodwill evolving solely based on what advice is communicated, this claim should stand no matter what feedback one thinks actors will receive ex-post about past interactions in such relationships (see Remark 2).

3.2. Senders with transparent motives

Because money burning and bonuses can expand the set of possible equilibrium payoffs for S, repetition can do the same. However, this is not so for every sort of expert. One strand of the cheap talk literature (e.g. Chakraborty and Harbaugh, 2010; Lipnowski and Ravid, 2020) has focused on the case in which experts have transparent motives—modeled by $u_S(a, \theta)$ being constant in its second argument. Pairing results from that literature with our theorem yields the following immediate consequence.

Corollary 2 (*Irrelevance under transparency*). *Suppose S's objective u_S is state-independent. Then $V_\delta = V_0$.*

The corollary follows from the securability theorem (Theorem 1 from Lipnowski and Ravid, 2020) in the theory of cheap talk games with state-independent S preferences. Say a payoff v is *securable* if S has some way of communicating such that a Bayesian R will, message by message, have some incentive-compatible action that gives S a payoff of at least v . The securability theorem says that, if a payoff v is securable and higher than the payoff S obtains in a babbling equilibrium, then some equilibrium of the cheap talk game delivers S a payoff of v . Loosely, if some messages are “too persuasive” and so give S a payoff strictly higher than v , communication can be made noisier in such a way that every message ensures some incentive-compatible R (mixed) action gives S a payoff of exactly v —which implies v is an equilibrium payoff for S because her preferences are state independent.

Let us observe that the securability theorem can be reinterpreted as a result about money burning, and so (given Theorem 1) can be applied to the repeated game. Indeed, suppose some S payoff v is attainable in an equilibrium of an auxiliary game with (capped) money burning. Because S has state-independent preferences, incentive compatibility means every on-path choice (of message and burned money) she makes must lead her to a payoff of exactly v . But then, whatever R behavior she induces gives her a payoff of *at least* v when we ignore the money she has burned. Hence, v is securable: the securability theorem says v is an equilibrium S payoff of the cheap talk game without money burning if it is higher than her babbling payoff. In light of Theorem 1, the payoff in $\bar{v}_\delta$ belongs to V_0 . An analogue of the securability theorem for S payoffs below her babbling payoff (Lipnowski and Ravid, 2020, footnote 15) can similarly be applied to the auxiliary game with bonuses, and so analogously tells us $\underline{v}_\delta \in V_0$; thus, $V_\delta = V_0$.

3.3. Repetition as commitment

An easy corollary of our theorem is a form of anti-folk theorem. Indeed, define the commitment (a.k.a. Bayesian persuasion) value,

$$v^{BP} := \max_{p \in \Delta \Delta \Theta: \int \mu \, dp(\mu) = \mu_0} \int_{\Delta \Theta} \max_{a^*(\mu) \in A} \int u_S(a^*(\mu), \cdot) \, d\mu \, dp(\mu) \\ \text{s.t. } a^*(\mu) \in \operatorname{argmax}_{a \in A} \int u_R(a, \cdot) \, d\mu$$

Since Kamenica and Gentzkow (2011), a large active literature has studied optimal information structures by an expert who can commit in order to persuade a decision maker; such an expert's optimal value is given by v^{BP} .¹⁵ A natural question is whether relational incentives suffice to generate commitment power.

Given that R myopically best responds in any equilibrium, it is immediate that the sender payoff is bounded above by v^{BP} in our model. But can this value be obtained? In a model with short-run receivers (which is strategically equivalent to our assumption of R being unmonitored) but perfect ex-post feedback about S's information, Best and Quigley (2020) fully settle this question when S has transparent motives (and the state space is finite): Repetition restores the Bayesian persuasion value to S if and only if v^{BP} can be attained with partitional information, that is, with S partitioning the set of states and reporting to which partition element the true state belongs.

The next corollary gives a general, definitive answer for our setting, with no economic assumptions on players' objectives. Specifically, it says that varying S's goodwill based solely on what advice she gives can never fully bridge a commitment gap. More formally, if all equilibria from the one-shot game with static cheap talk are strictly worse for S than the Bayesian persuasion value, then all equilibria from the repeated cheap talk game are strictly worse for S than the Bayesian persuasion value as well.

Corollary 3 (Anti-folk theorem). *If $\bar{v}_0 < v^{BP}$, then $\sup_{\delta \in [0,1)} \bar{v}_\delta < v^{BP}$ too.*

¹⁵ Prior work in the theory of undiscounted repeated zero-sum games (Aumann et al., 1995; Aumann and Maschler, 1966) has shown that such commitment power can naturally arise in long-run relationships—although that work proceeds with a perfectly persistent state.

This result is immediate from the first corollary. Indeed, consider again an S-preferred equilibrium of the game of cheap talk with money burning. There are two cases. First, if no money is burned on path, then the theorem implies that repetition does not raise S's highest equilibrium value relative to one-shot cheap talk. If, conversely, money is burned on path, then the expected amount of money burned is a lower bound on the residual commitment gap. Indeed, if S could commit to a communication rule, she could always provide exactly the same information as in the money-burning game, without incurring the cost of burned money.

We should emphasize that the above corollary relies (as do all of our results) on our specific monitoring assumptions. Rather than demonstrating that relational incentives can never sustain commitment power in communication, then, Corollary 3 serves as a useful theoretical benchmark. Whereas in general the ability to use feedback to sustain commitment power depends on specific details of the stage game—as evidenced by Best and Quigley's (2020) results for the case of perfect state feedback—our result shows that relational incentives can never, absent some supporting feedback, hope to replicate complete commitment power.

4. Application: project implementation

In this section, we apply our theorem to study a repeated project implementation game.

Environment. A firm's CEO (R) decides every period at what scale $a_t \in [0, 1]$ the firm will implement a project in the division of a manager (S). In line with our model, the CEO's choice of a_t is private and not observed by the manager. Each project produces a marginal expected revenue of θ_t to the firm, but also entails a constant marginal cost c of implementation. The firm's flow profit is therefore $a_t(\theta_t - c)$. Project values θ are atomlessly distributed with support $[0, 1]$, and the cost parameter satisfies $\mathbb{E}[\theta] < c < 1$. The manager is driven by empire-building motives, and so internalizes only the benefits $a_t\theta_t$. So, while the CEO must rely on the manager's expertise, a conflict of interest precludes perfect communication. Importantly, we assume that the project's realized return is extremely noisy and/or delayed, so that providing incentives to the manager based on ex-post outcomes is impractical.¹⁶

This example is a continuous version of the 2×2 model of Lipnowski and Ramos (2020), with the crucial difference that the CEO's implementation choices are private here, and perfectly observable in that paper. Given the substantially different observability assumptions, neither their techniques nor their results can be directly applied here. However, as we show below, despite the extreme paucity of instruments, a repeated interaction still facilitates some communication. Applying our tools, we quantify the gains of repeated interaction for a given discount factor; we proceed informally for brevity.

As a starting observation, notice that babbling is the unique equilibrium of the stage game, giving both players a payoff of 0 (their minimax payoffs). Indeed, no two on-path messages could lead to different expected implementation scales, for the lower one could not be incentive-compatible for any nonzero S type. But then, R rationally chooses $a = 0$ almost surely, because $\mathbb{E}\theta < c$. This tells us that the equilibrium S payoff set is always of the form $[0, \bar{v}_\delta]$ and, in particular, bonuses cannot hurt S in this example.

¹⁶ Recall, the assumptions of our game require that players not observe their own payoffs. Rather than interpreting such an assumption literally, we take the view that a firm's choices and their exact profitability might, sometimes, be prohibitively difficult to assess on a time horizon short enough to provide meaningful dynamic incentives.

Pricing advice with burned money. Given the theorem, we begin by analyzing the game of one-shot cheap talk with M -capped money burning, for some cap $M \in \mathbb{R}_+$. Toward finding S 's best equilibrium payoff for this auxiliary game, we first invest in some notation. First, observe that there is a unique $\theta_* \in (0, c)$ such that $\mathbb{E}[\theta \mid \theta \geq \theta_*] = c$.¹⁷ In what follows, let $\bar{v} := \mathbb{E}[(\theta - \theta_*)_+] > 0$ and $\bar{\delta} := \frac{\theta_*}{\mathbb{E}[\max\{\theta, \theta_*\}]} = \frac{\theta_*}{\theta_* + \bar{v}} \in (0, 1)$.

In contrast to the case without money burning, information can be communicated effectively by using burned money to price different recommendations. For example, it would be natural to look for a *cutoff equilibrium* where all types below an interior cutoff burn no money and all types above the same burn a fixed amount $m \in (0, M]$ of money, and where no further information is communicated. The CEO, being even more pessimistic than at the prior belief, would respond to no burned money by choosing action zero. Clearly, for the CEO to willingly choose a different action following m , it must be that the cutoff is at least θ_* . Let us focus on the case that the cutoff is exactly θ_* , so that any action $\hat{a} \in [0, 1]$ is optimal for the CEO when the manager burns m . Next, given single crossing of the manager's preferences, she will find this play optimal if and only if the cutoff type is indifferent. That is, this cutoff play is an equilibrium if and only if $\hat{a}\theta_* = m$, i.e. recommendation $\hat{a}$ is priced (at m) exactly fairly from type θ_* 's perspective.

Now, what is the best equilibrium within this cutoff class?

Claim 1. *The game with M -capped money burning admits cutoff equilibria with cutoff θ_* . In the best such equilibrium for the manager, the highest on-path implementation scale is $\min\{1, \frac{M}{\theta_*}\}$, and generates a manager payoff of $\min\{\bar{v}, \frac{M\bar{v}}{\theta_*}\}$.*

Given the described behavior by the CEO, the informed manager chooses optimally between burning nothing and paying a price of m , so that her expected payoff is $\mathbb{E} \max\{0, \hat{a}\theta - m\} = \hat{a}\bar{v}$, where $\hat{a}$ is the CEO's chosen implementation scale when m is burned. Therefore, she benefits from having $\hat{a} \in [0, 1]$ as large as possible subject to the constraint that $\hat{a}\theta_* = m \leq M$. And indeed, the best such equilibrium has $\hat{a} = \min\{1, \frac{M}{\theta_*}\}$, and so generates a payoff of $\hat{a}\bar{v} = \min\{\bar{v}, \frac{M\bar{v}}{\theta_*}\}$.

Best equilibrium with burned money. Next, we argue that the cutoff equilibrium described above is in fact a best equilibrium for the manager. To this end, we take a mechanism design approach. Consider an arbitrary equilibrium, and for each $\theta \in [0, 1]$, let $\alpha(\theta)$, $\tau(\theta)$, and $u(\theta)$ denote that type's interim expected implementation scale, burned money, and equilibrium utility, respectively. So $u(\theta) = \theta\alpha(\theta) - \tau(\theta)$; and, if type θ were to play some type $\tilde{\theta}$'s equilibrium strategy, she would get a lower payoff of $\theta\alpha(\tilde{\theta}) - \tau(\tilde{\theta})$. As manager incentives imply that no type strictly prefers to use another type's mixture over $[0, M] \times Z$, we can employ the standard toolbox of one-dimensional mechanism design (Myerson, 1981). In particular, we learn that α is nondecreasing and that $u(\theta) - u(0) = \int_0^\theta \alpha$ for every type θ , where $\int_0^\theta \alpha$ is a short-hand for $\int_0^\theta \alpha(\tilde{\theta}) d\tilde{\theta}$. Some features of the allocation follow directly from these incentive conditions. For instance:

Claim 2. *In any equilibrium of the game with M -capped money burning, no project of type in $[0, \theta_*)$ is implemented at a strictly positive scale.*

¹⁷ From the analysis of the salesperson example in Kamenica and Gentzkow (2011), it is straightforward to show that revealing whether or not the state is above θ_* would be uniquely manager-optimal under commitment, leaving the CEO with his babbling payoff.




Fig. 1. A cutoff equilibrium dominates another equilibrium: Interim utilities agree to the left of the cutoff, and then grow more quickly for the cutoff equilibrium. Hence, the latter generates higher S values.

To verify the claim, let us focus on the nontrivial case that α is not globally zero; let $\hat{\theta} \in [0, 1]$ be the infimum type at which $\alpha > 0$. By strict single crossing of the manager's preferences, every type $\theta > \hat{\theta}$ is certain to choose a message-money pair in Ω_1 , the set of all $(y, z) \in [0, M] \times Z$ that lead to a strictly positive expected implementation scale; and every type $\theta < \hat{\theta}$ is certain to choose a message-money pair outside of Ω_1 . Moreover, linearity of payoffs in the implementation scale implies that the CEO would find it optimal to implement the project at full scale whenever a signal in Ω_1 is sent, and at zero scale otherwise. That the profit $\mathbb{E}[(\theta - c)\mathbf{1}_{\theta > \hat{\theta}}] = \mathbb{E}[(\theta - c)\mathbf{1}_{\theta \geq \hat{\theta}}]$ this CEO behavior would generate is nonnegative tells us that $\hat{\theta} \geq \theta_*$, as desired.

We can now construct a modified equilibrium of the cutoff form which raises the manager's payoff. Her indirect utility under the original and new equilibria are depicted in Fig. 1.

Because $\alpha(\theta) = 0$ for $\theta \leq \theta_*$, and from monotonicity of α , the expected implementation scale of a project of type θ is bounded above by $\tilde{\alpha}(\theta) := \alpha(1)\mathbf{1}_{\theta \geq \theta_*}$. The manager's interim equilibrium payoff, when the project is type θ , is then

$$u(\theta) = u(0) + \int_0^\theta \alpha = -\tau(0) + \int_0^\theta \alpha \leq 0 + \int_0^\theta \tilde{\alpha} = \alpha(1)(\theta - \theta_*)_+.$$

Therefore, the manager is better off with a cutoff equilibrium whose high implementation scale is $\hat{a} = \alpha(1)$. All that remains is to observe that the implied burned money m for high types in such an equilibrium does not exceed the cap. And, indeed,

$$m = \alpha(1)\theta_* = \alpha(1) - \int_{\theta_*}^1 \alpha(1) \leq \alpha(1) - \int_{\theta_*}^1 \alpha = \alpha(1) - u(1) + u(\theta_*) = \tau(1) + 0 \leq M,$$

where the first inequality follows from α being monotone, and the second inequality follows from the original equilibrium respecting the money burning cap of M . So the modified cutoff equilibrium is feasible. Summarizing the above arguments yields:

Claim 3. Any equilibrium of the game with M -capped money burning admits a cutoff equilibrium with cutoff θ_* that is weakly better for the manager.

Combined, the above analysis shows that a best equilibrium for the manager with money burning capped at M takes a simple form, and generates an ex-ante payoff of $\min\{\bar{v}, \frac{M\bar{v}}{\theta_*}\}$ to the manager.

Equilibrium manager payoffs in the repeated game. Having characterized the manager's attainable payoffs with capped money burning, we are poised to apply the theorem to our project implementation example.

Proposition 1. V_δ is equal to $\{0\}$ if $\delta < \bar{\delta}$, and $[0, \bar{v}]$ if $\delta \geq \bar{\delta}$.

We have already established $V_\delta = [0, \bar{v}_\delta]$. Given the theorem, $\bar{v}_\delta$ is the largest number $v \geq 0$ such that v is less than or equal to the payoff $\min\{\bar{v}, \frac{M\bar{v}}{\theta_*}\}$, where $M = \frac{\delta}{1-\delta}v$. Rephrasing, $\bar{v}_\delta$ is the largest $v \in [0, \bar{v}]$ such that $v \leq \frac{\delta\bar{v}}{(1-\delta)\theta_*}v$. The latter inequality simply requires that v be non-negative if $\frac{\delta\bar{v}}{(1-\delta)\theta_*} > 1$ and nonpositive if $\frac{\delta\bar{v}}{(1-\delta)\theta_*} < 1$. Finally, observe that the expression $\frac{\delta\bar{v}}{(1-\delta)\theta_*}$ is strictly increasing in δ , and equal to 1 when δ is equal to $\bar{\delta}$. The proposition's characterization of $\bar{v}_\delta$ follows directly.

A manager-optimal equilibrium. Let us now construct an equilibrium that attains $\bar{v}$ when $\delta \geq \bar{\delta}$: it follows directly from the proposition that it is a best equilibrium for the manager.¹⁸ Let $\varepsilon := \frac{\delta - \bar{\delta}}{\delta(1-\bar{\delta})} \in [0, 1)$, and fix two distinct messages $z_H, z_L \in Z$. The manager has a history-independent reporting strategy in which she says z_H whenever the current state is at least θ_* , and says z_L otherwise. The CEO has a history-contingent strategy in which he implements a project at scale ε^{k-1} the k^{th} time the manager says z_H (for any $k \in \mathbb{N}$), and implements no project (i.e. $a = 0$) whenever anything else is reported. Direct computation (see Appendix A.3) shows that, paired with an appropriate belief map, this pure strategy profile yields an equilibrium that gives the manager a payoff of $\bar{v}$.

This equilibrium demonstrates how the manager's future goodwill is used to price today's recommendations. As the relationship progresses and her goodwill is gradually exhausted, the same advice gives a smaller benefit, but the cost—fraction $1 - \varepsilon$ of her residual goodwill—shrinks as well.

Appendix A

This appendix comprises four subsections. First, we provide a useful alternative formalism for reasoning about S payoffs in communication games. Second, we prove formally our main theorem—which is entirely standard given the background work in the first subsection. Third, we provide supporting computations for the S-optimal equilibrium we construct in our project implementation application. Fourth, we demonstrate by example that bonuses are, in general, a necessary ingredient of our equilibrium characterization result.

¹⁸ The construction here, which exploits some linearity in the application, does not specialize the general-purpose construction in the proof of our main theorem, as described at the end of Section 2. Whereas here we see the manager's goodwill being gradually depleted as she asks for projects, the general construction would instead see it stochastically disappearing.

A.1. Equilibrium payoffs of static communication games

In this subsection, we develop a useful language to reason about equilibria of our auxiliary one-shot games. In short, we will identify an equilibrium with the collection of money-action-belief triples that can be induced on-path, analogous to the equivalence between direct mechanisms and delegated sets in the literature on optimal delegation (e.g. Holmström, 1982).

In what follows, we let $\bar{M} \in \mathbb{R}_+$ be high enough that any payoff attainable with $\bar{M}$ -capped bonuses or money burning for some $M \geq 0$ is also attainable with some cap $\tilde{M} \leq \bar{M}$ —for instance as defined in Footnote 11. We next let $X := [-\bar{M}, \bar{M}] \times \Delta A \times \Delta \Theta$, and let $\mathbb{K}X$ be the set of closed nonempty subsets of X , endowed with the Hausdorff metric. Both are compact metrizable spaces because A and Θ are.

Loosely, we will identify an equilibrium with a subset of X . In slightly more detail: each signal (i.e. money-message pair) sent in equilibrium will induce a monetary consequence (y), a mixed action (α), and a belief (μ)—and so induces an element of X . But then, the equilibrium distribution of messages yields a distribution over X . We will observe that the support of such a distribution is enough to compute an S payoff, and that it is possible to directly characterize the set of all subsets of X that can arise as such a support.

A few notations will prove useful for our analysis.

Notation 1. For any $K \subseteq X$, define the projections

$$K_1 := \{y : (y, \alpha, \mu) \in K\}, \quad K_3 := \{\mu : (y, \alpha, \mu) \in K\}, \quad K_{32} := \{(\mu, \alpha) : (y, \alpha, \mu) \in K\}.$$

Below, we define several more notations that will be useful. Informally, their descriptions are as follows. For each R belief μ , $A^*(\mu)$ gives his set of (mixed) best responses. For each “signal” $x \in X$ and S type θ , the value $f(x, \theta)$ is S’s interim value from sending that signal. For each “signal set” $K \subseteq X$, the value $f^*(K, \theta)$ gives S type θ ’s interim value from facing that signal set and choosing optimally, and so $f^{**}(K)$ is (as we will show) the associated greatest deviation gain of any S type. The set $\mathcal{K}$ is the set of all “signal sets” consistent with R being Bayesian and both players being rational, i.e. with every feature of equilibrium except for the bounds on the monetary signals S uses.¹⁹ The set $\mathcal{Y}$ is a set of sets of monetary levels that includes all relevant restrictions on monetary signals. For any restriction Y on monetary signals, $\Gamma(Y)$ is the associated set of equilibrium “signal sets”, $\mathcal{W}(Y)$ is the associated set of equilibrium interim S payoff functions, and $W(Y)$ is the associated set of equilibrium ex-ante S payoffs.

Having interpreted these objects, we define them formally.

Definition 3. Define the following objects:

$$\begin{aligned} A^* : \Delta \Theta &\rightrightarrows \Delta A \\ \mu &\mapsto \Delta \left[\operatorname{argmax}_{a \in A} \int_{\Theta} u_R(a, \cdot) \, d\mu \right] \\ f : X \times \Theta &\rightarrow \mathbb{R} \end{aligned}$$

¹⁹ When specialized to the case of finitely many states and actions, one can verify that $W(\{0\})$ is exactly the set of sender payoffs derived by Aumann and Hart (2003) in their one-shot analysis, evaluated at μ_0 .

$$(y, \alpha, \mu, \theta) \mapsto u_S(\alpha, \theta) - y$$

$$f^* : \mathbb{K}X \times \Theta \rightarrow \mathbb{R}$$

$$(K, \theta) \mapsto \max_{x \in K} f(x, \theta)$$

$$f^{**} : \mathbb{K}X \rightarrow \mathbb{R}_+$$

$$K \mapsto \max_{x=(y, \alpha, \mu) \in K} \int_{\Theta} [f^*(K, \cdot) - f(x, \cdot)] d\mu$$

$$\mathcal{K} := \left\{ K \in \mathbb{K}X : \underbrace{\bar{\text{co}}(K_3) \ni \mu_0}_{\text{Bayes}}, \underbrace{K_{32} \subseteq \text{gr}(A^*)}_{IC_R}, \underbrace{f^{**}(K) = 0}_{IC_S} \right\}$$

$$\mathcal{Y} := \mathbb{K}[-\bar{M}, \bar{M}]$$

$$\Gamma : \mathcal{Y} \rightrightarrows \mathcal{K}$$

$$Y \mapsto \{K \in \mathcal{K} : K_1 \subseteq Y\}$$

$$\mathcal{W} : \mathcal{Y} \rightrightarrows C(\Theta)$$

$$Y \mapsto \{f^*(K, \cdot) : K \in \Gamma(Y)\}$$

$$W : \mathcal{Y} \rightrightarrows \mathbb{R}$$

$$Y \mapsto \left\{ \int_{\Theta} w d\mu_0 : w \in \mathcal{W}(Y) \right\}.$$

Let us first document some basic topological properties for the above objects.

Lemma 1. A^* is upper hemicontinuous²⁰; f, f^*, f^{**} are continuous; $\mathcal{K}$ is closed; and $\Gamma, \mathcal{W}, W$ are nonempty-compact-valued, upper hemicontinuous, and monotone with respect to set inclusion.

Proof. First, f is obviously continuous. Next, A^* is upper hemicontinuous and f^* is continuous by Berge's theorem (note that the correspondence $\mathbb{K}X \rightrightarrows X$ taking $K \mapsto K$ is a continuous correspondence), implying f^{**} is continuous by the same argument. As f^{**} is continuous, A^* is upper hemicontinuous, and the $\bar{\text{co}}$ operator and projection operators are continuous on $\mathbb{K}X$, it follows that $\mathcal{K}$ is closed in $\mathbb{K}X$.

Let us observe now that the graph of Γ is a closed subset of $\mathcal{Y} \times \mathcal{K}$ (making it compact-valued and upper hemicontinuous). Consider a sequence $\{Y^n\}_n \subseteq \mathcal{Y}$ converging to Y and $K^n \in \Gamma(\mathcal{Y}_n)$

²⁰ In the remaining argument, R only enters the analysis through A^* , and the only features of A^* that we use are nonempty-compact-valuedness and upper hemicontinuity. Accordingly, our main theorem would apply without change if there were multiple receivers.

for each n with $K^n \rightarrow K \in \mathbb{K}X$. As $\mathcal{K}$ is closed, we know $K \in \mathcal{K}$; and continuity of projection yields $K_1^n \rightarrow K_1$. Finally, that $K_1^n \subseteq Y^n$ for each n implies that their limit is a subset of Y . Therefore, $K \in \Gamma(Y)$, as required.

Now, because f^* is continuous, and so uniformly continuous on its compact domain, it follows that the map $\Phi : \mathbb{K}X \rightarrow C(\Theta)$ given by $\Phi(K) := f^*(K, \cdot)$ is continuous. But then $\mathcal{W}$ is a continuous transformation of Γ , and so too is its continuous transformation W . It follows that (since Γ is) both $\mathcal{W}$ and W are compact-valued and upper hemicontinuous.

Finally, it is obvious that $\Gamma, \mathcal{W}, W$ are monotone with respect to set inclusion. That they are nonempty-valued is then easy to show by observing $\{(y, \alpha, \mu_0)\} \in \Gamma(Y)$ for any $y \in Y \in \mathcal{Y}$ and $\alpha \in A^*(\mu_0)$. ■

With these notations and preliminary observations in hand, we are equipped to complete our characterization. The following lemma shows that attainable S interim payoff vectors can be fully described by closed subsets of X they induce (and so are exactly given by $\mathcal{W}$). Moreover, it shows that it is without loss of generality for S payoffs to focus on equilibria in which the monetary signal chosen in equilibrium is a deterministic function of the message sent. This latter feature will prove useful in the proof of our main equivalence theorem because S 's continuation value (which will serve as a stand-in for her monetary signal) is a function of her current message rather than a separate choice—a difference that the equivalence between (i) and (ii) renders immaterial.

Lemma 2. *Given convex $Y = [y, \bar{y}] \in \mathcal{Y}$ and $w(\cdot) \in \mathbb{R}^\Theta$, the following are equivalent:*

- (i) *There exists an equilibrium $\langle \sigma, \rho, \beta \rangle$ of the game with cheap talk and $[y^*, \bar{y}]$ -monetary signals yielding S interim payoff function w , for some $y^* \in Y$.*
- (ii) *There exists $\langle \sigma, \rho, \beta \rangle$ as in (i) such that $\sigma(\cdot | \theta) = \hat{\sigma}(\cdot | \theta) \circ (\hat{y}, id_Z)^{-1}$ for all $\theta \in \Theta$ for some measurable $\hat{\sigma} : \Theta \rightarrow \Delta Z$ and $\hat{y} : Z \rightarrow Y$.*
- (iii) *$w \in \mathcal{W}(Y)$.*

Proof. As $(ii) \implies (i)$ obviously, we prove that $(i) \implies (iii) \implies (ii)$.

- $(i) \implies (iii)$: Letting $\langle \sigma, \rho, \beta \rangle$ be an equilibrium as in (i), we must find $K \in \Gamma(Y)$ such that $w = f^*(K, \cdot)$.

Let $\mathbb{P} \in \Delta X$ be given by, for all Borel $\hat{Y} \subseteq Y$, $\hat{A} \subseteq \Delta A$, $\hat{B} \subseteq \Delta \Theta$,

$$\mathbb{P}(\hat{Y} \times \hat{A} \times \hat{B}) := \int_{\Theta} \sigma \left([\hat{Y} \times Z] \cap \rho^{-1}(\hat{A}) \cap \beta^{-1}(\hat{B}) \middle| \cdot \right) d\mu_0.$$

Then let $K := \text{supp}(\mathbb{P})$. By Bayes, $\mu_0 = \int_{\Delta \Theta} \mu \, d \text{marg}_{\Delta \Theta} \mathbb{P}(\mu) \in \overline{\text{co}}(K_3)$. IC_R then implies (β, ρ) is $\text{gr}(A^*)$ -valued, so that $\alpha \in A^*(\mu)$ a.s.- $\mathbb{P}(y, \alpha, \mu)$. But then, that A^* is upper hemicontinuous means $\alpha \in A^*(\mu) \, \forall (\mu, \alpha) \in K_{32}$.

Assume for a contradiction that $f^{**}(K) \neq 0$. Then there exists $x = (y, \alpha, \mu) \in K$ such that $\mu\{f(x, \cdot) < f^*(K, \cdot)\} > 0$. Continuity of f and f^* delivers a nonempty open $\check{X} \subseteq X$ of x such that $\check{\mu}\{f(\check{x}, \cdot) < f^*(K, \cdot)\}$ for each $\check{x} = (\check{y}, \check{\alpha}, \check{\mu}) \in \check{X}$. But, as the supremum value of a given type $\theta \in \Theta$ is at least $f^*(K, \theta)$, IC_S then implies that no $\theta \in \bigcup_{\check{x} \in \check{X}} \{f(\check{x}, \cdot) < f^*(K, \cdot)\}$ sends a signal inducing an element of $\check{X}$ in equilibrium.²¹ This contradicts Bayes

²¹ By a signal inducing an element of $\check{X}$, we mean $(\check{y}, \check{z}) \in Y \times Z$ with $(\check{y}, \rho(\check{y}, \check{z}), \beta(\check{y}, \check{z})) \in \check{X}$.

(given that messages corresponding to $\check{X}$ must be positive probability by definition of K). So $f^{**}(K) = 0$, and therefore $K \in \Gamma(Y)$.

The equilibrium interim payoff vector is then:

$$\sup_{(y,z) \in Y \times Z} [u_S(\rho(y, z), \cdot) - y] = f^*(K, \cdot) \in \mathcal{W}(Y).$$

- (iii) $\implies$ (ii): Suppose $K \in \Gamma(Y)$. Fix some $(y^*, \alpha^*, \mu^*) \in \operatorname{argmin}_{(y, \alpha, \mu) \in K} y$, and let $Y^* := [y^*, \bar{y}]$. For convenience, we assume without loss (appealing to Kuratowski's theorem since Z is an uncountable Polish space) that $Z = \Delta\Theta$. A measurable selection theorem (Aliprantis and Border, 2006, Theorem 18.13) delivers some measurable $(\hat{y}, \hat{\alpha}) : K_3 \rightarrow K_{12}$ such that $(\hat{y}(\mu), \hat{\alpha}(\mu), \mu) \in K$ for all $\mu \in K_3$ and $(\hat{y}(\mu^*), \hat{\alpha}(\mu^*)) = (y^*, \alpha^*)$. Next, that $\mu_0 \in \overline{\operatorname{co}}(K_3)$ tells us (see Phelps, 2001) that there is some $p \in \Delta(K_3)$ such that $\int_{K_3} \mu \, dp(\mu) = \mu_0$. As is now standard (e.g. Kamenica and Gentzkow, 2011), there exists some measurable $\hat{\sigma} : \Theta \rightarrow \Delta(K_3)$ such that $\int_{\Theta} \hat{\sigma} \, d\mu_0 = p$ and, for all Borel $\hat{\Theta} \subseteq \Theta$, $\hat{B} \subseteq K_3$, we have that

$$\int_{\hat{\Theta}} \hat{\sigma}(\hat{B} \mid \cdot) \, d\mu_0 = \int_{\hat{B}} \mu(\hat{\Theta}) \, d\rho(\mu).$$

Now, define

$$\begin{aligned} \sigma : \Theta &\rightarrow \Delta([y, \bar{y}] \times Z) \\ \theta &\mapsto \hat{\sigma}(\cdot \mid \theta) \circ (\hat{y}, \operatorname{id}_Z)^{-1} \end{aligned}$$

and

$$\begin{aligned} (\beta, \rho) : [y^*, \bar{y}] \times Z &\rightarrow \Delta\Theta \times \Delta A \\ (y, z) &\mapsto \begin{cases} (\mu, \hat{\alpha}(\beta)) & : z \in K_3 \text{ and } y = \hat{y}(z) \\ (\mu^*, \alpha^*) & \text{otherwise.} \end{cases} \end{aligned}$$

By construction, this is an equilibrium of the game with cheap talk and $[y^*, \bar{y}]$ -monetary signals (of the form in (ii)) generating payoff vector $f^*(K, \cdot)$ to S . ■

A.2. Equilibrium payoffs of repeated communication games

Given the results of the previous section, we can now define, and prove essential properties of, the present analogue of APS's set operator.

Notation 2. Define:

$$\begin{aligned} \bar{g} : \bar{\mathbb{R}}_+ &\rightarrow \mathbb{R} \\ M &\mapsto \max W([0, \bar{M} \wedge M]), \\ \underline{g} : \bar{\mathbb{R}}_+ &\rightarrow \mathbb{R} \\ M &\mapsto \min W([-(\bar{M} \wedge M), 0]), \\ G : \bar{\mathbb{R}}_+ &\rightrightarrows \mathbb{R} \\ M &\mapsto [\underline{g}(M), \bar{g}(M)]. \end{aligned}$$

Lemma 3. *G is a nonempty-compact-valued, upper hemicontinuous, and increasing with respect to set inclusion.*

Proof. We invoke Lemma 1. First, W is increasing with respect to set inclusion. This implies that $\bar{g}$ is increasing, $\underline{g}$ is decreasing, and $\emptyset \neq W(\{0\}) \subseteq G(M)$ for all $M \in \bar{\mathbb{R}}_+$. Therefore, G is increasing and nonempty-valued; it is compact-valued by definition. Next, W is upper hemicontinuous and compact-valued, which implies that $\bar{g}$ is upper semicontinuous, $\underline{g}$ is lower semicontinuous, and both are bounded. From this, we conclude that G is upper hemicontinuous and compact-valued. ■

Lemma 4. *Given $\langle \sigma, \beta, \rho \rangle$, R rationality holds if and only if R myopically best responds to per-period beliefs.*

The above lemma is immediate from the fact that R 's behavior is not (even imperfectly) observed by S .

Now, given that states are independent of past states and actions, and R is not monitored, we observe that public strategies (which enjoy recursive structure) are without loss.

Definition 4. Given a time- t history h^i for player $i \in S, R$, let h^P denote its projection onto $([0, 1] \times Z)^t \times [0, 1]$, i.e. the associated public history.

Say $\langle \sigma, \beta, \rho \rangle$ is a **public equilibrium** if it is an equilibrium such that, for all $h^S, \tilde{h}^S \in \mathcal{H}^S$ [resp. $h^R, \tilde{h}^R \in \mathcal{H}^R$] such that $h^P = \tilde{h}^P$, we have $\sigma(\cdot|h^S, \cdot) = \sigma(\cdot|\tilde{h}^S, \cdot)$ [resp. $\rho(\cdot|h^R, \cdot) = \rho(\cdot|\tilde{h}^R, \cdot)$ and $\beta(\cdot|h^R, \cdot) = \beta(\cdot|\tilde{h}^R, \cdot)$].

Lemma 5. *For every $v \in V_\delta$, there is a public equilibrium generating S payoff v .*

We omit the proof of the above lemma, which is a standard adaptation of the usual constructive argument for games with product monitoring structure (e.g. Fudenberg and Levine, 1994, Theorem 5.2); see (Mailath and Samuelson, 2006, Proposition 10.1.1) for further discussion. As players' private histories are independent conditional on the public history, for any conditioning a player might do in some equilibrium on private past information, there is a payoff-equivalent (to both players, even if the other player were to unilaterally deviate) public strategy in which current private contemporaneous randomization replaces this conditional play.

Lemma 6. *Let $\bar{v} := \sup V_\delta$, $\underline{v} := \inf V_\delta$, and $M := \frac{\delta}{1-\delta}(\bar{v} - \underline{v})$. Then, $V_\delta \subseteq G(M)$.*

Proof. The claim is immediate when $\delta = 0$, so focus is on the case of $\delta \in (0, 1)$. Consider any public (which is without loss of payoffs by Lemma 5) equilibrium without initial public randomization; say it generates sender value $v^* \in V_\delta$. Let initial play be given by $\sigma_0 : \Theta \rightarrow \Delta Z$, $\rho_0 : Z \rightarrow \Delta A$, and let $v'(z) \in V_\delta$ be the sender's continuation value following message z , $\forall z \in Z$. Now $\forall z \in Z$, let

$$y(z) := \frac{\delta}{1-\delta}[v^* - v'(z)] \subseteq \left[\frac{\delta}{1-\delta}(v^* - \bar{v}), \frac{\delta}{1-\delta}(v^* - \underline{v}) \right].$$

Let $\mathbb{P} \in \Delta X$ be given by, for all Borel $\hat{Y} \subseteq [-\bar{M}, \bar{M}]$, $\hat{A} \subseteq \Delta A$, $\hat{B} \subseteq \Delta \Theta$,

$$\mathbb{P}(\hat{Y} \times \hat{A} \times \hat{B}) := \int_{\Theta} \sigma \left([v']^{-1} \left(\frac{1-\delta}{\delta} \hat{Y} + v^* \right) \cap \rho^{-1}(\hat{A}) \cap \beta^{-1}(\hat{B}) \right) d\mu_0.$$

Let $K := \text{supp}(\mathbb{P})$. It is straightforward that $\mu_0 \in \overline{\text{co}}(K_3)$ by Bayes, that $K_{32} \subseteq \text{gr}(A^*)$ by IC_R given Lemma 4, and that $f^{**}(K) = 0$ by IC_S (specifically, that S has no profitable one-shot deviation). Therefore, $K \in \mathcal{K}$. It follows that

$$K \in \Gamma \left(\left[\frac{\delta}{1-\delta}(v^* - \bar{v}), \frac{\delta}{1-\delta}(v^* - \underline{v}) \right] \right),$$

and so

$$v^* \in W \left(\left[\frac{\delta}{1-\delta}(v^* - \bar{v}), \frac{\delta}{1-\delta}(v^* - \underline{v}) \right] \right).$$

The above applies to every v^* that is attainable in equilibrium without initial public randomization. But such equilibrium values can clearly approximate each of $\bar{v}$ and $\underline{v}$ arbitrarily well. So, applying upper hemicontinuity of W and taking limits as $v^* \rightarrow \bar{v}$ and as $v^* \rightarrow \underline{v}$, we get that $\bar{v} \in W([0, M])$ and $\underline{v} \in W([-M, 0])$, i.e. $[\underline{v}, \bar{v}] \subseteq G(M)$. ■

Lemma 7. Suppose that $\underline{v} \leq \bar{v}$ and $M = \frac{\delta}{1-\delta}(\bar{v} - \underline{v})$ are such that $[\underline{v}, \bar{v}] \subseteq G(M)$. Then $[\underline{v}, \bar{v}]$ is a subset of V_δ .

Proof. Let $v_0 := g(M)$ and $v_1 := \bar{g}(M)$. By hypothesis, $v_0 \leq \underline{v} \leq \bar{v} \leq v_1$. We will show by construction that $\{v_0, v_1\}$ are both S equilibrium values, and the lemma will then follow from public randomization. Focus on the nontrivial case that $v_0 < v_1$ (since repetition of a babbling equilibrium yields the desired payoff in the complementary case).

Let us define a two-state automaton profile with states $\xi \in \{0, 1\}$. For $\xi \in \{0, 1\}$, Lemma 2 delivers an equilibrium $\langle \sigma_\xi, \rho_\xi, \beta_\xi \rangle$ of the game with cheap talk and Y_ξ -monetary signals, and measurable $\hat{\sigma}_\xi : \Theta \rightarrow \Delta Z$, $\hat{y}_\xi : Z \rightarrow Y_\xi$, delivering S payoff v_ξ , and such that

$$\sigma_\xi(\cdot | \theta) = \hat{\sigma}_\xi(\cdot | \theta) \circ (\hat{y}_\xi, \text{id}_Z)^{-1} \quad \forall \theta \in \Theta,$$

where $Y_1 = [0, M]$ and $Y_0 = [-\tilde{M}, 0]$ for some $\tilde{M} \in [0, M]$.

When the automaton is in the state $\xi \in \{0, 1\}$:

- Current play is given by $\langle \hat{\sigma}_\xi, \hat{\rho}_\xi, \hat{\beta}_\xi \rangle$, where

$$\hat{\rho}_\xi := \rho_\xi \circ (\hat{y}_\xi, \text{id}_Z) : Z \rightarrow \Delta A \text{ and } \hat{\beta}_\xi := \beta_\xi \circ (\hat{y}_\xi, \text{id}_Z) : Z \rightarrow \Delta \Theta.$$

- The automaton switches states if and only if

$$\underbrace{\omega_t}_{\text{public randomization} \sim \mathcal{U}[0,1]} < \frac{1-\delta}{\delta} \frac{|\hat{y}_\xi(z_t)|}{v_1 - v_0}.$$

Note, that $|\hat{y}_\xi| \leq M$ implies the switching cutoff is in $[0, 1]$.

A direct computation (see below) shows that the automaton yields S payoff of v_ξ when starting in state ξ , and incentive constraints from the auxiliary contracting game imply that no player will have a profitable one-shot deviation.²² Therefore, $\langle \hat{\sigma}_\xi, \hat{\rho}_\xi, \hat{\beta}_\xi \rangle_{\xi \in \{0,1\}}$ yields an equilibrium as required.

²² More specifically, an S type's deviation gain from sending message $\hat{z}$ instead of z for one period in state ξ will be proportional to her deviation gain from sending signal $(\hat{y}_\xi(\hat{z}), \hat{z})$ instead of $(\hat{y}_\xi(z), z)$ in the game with monetary signals.

Now, let us verify the construction yields an equilibrium that generates a payoff v_ξ in each state $\xi \in \{0, 1\}$. First, the Bayesian property and R incentive property are trivially inherited from the corresponding properties in the auxiliary games. Next, let us verify the S incentive property assuming the given strategy profile generates a payoff v_ξ from each initial state $\xi \in \{0, 1\}$. Appealing to the one-shot deviation principal, we need only show, for every $\xi \in \{0, 1\}$ and $\theta \in \Theta$, that S does not want to engage in a one-shot deviation and send a different report than the putative equilibrium prescribes. To see this, observe that her continuation payoff from sending message $z \in Z$ is

$$\begin{aligned}
 & (1 - \delta) \int_A u_S(\cdot, \theta) d\hat{\rho}_\xi(\cdot|z) d\hat{\sigma}_\xi(z|\theta) + \delta \left[v_\xi + (v_{1-\xi} - v_\xi) \frac{1-\delta}{\delta} \frac{|\hat{y}_\xi(z)|}{v_1 - v_0} \right] \\
 &= (1 - \delta) \int_A u_S(\cdot, \theta) d\hat{\rho}_\xi(\cdot|z) + \delta \left[v_\xi - (v_1 - v_0) \frac{1-\delta}{\delta} \frac{\hat{y}_\xi(z)}{v_1 - v_0} \right] \\
 &= \delta v_\xi + (1 - \delta) \left[\int_A u_S(\cdot, \theta) d\hat{\rho}_\xi(\cdot|z) - \hat{y}_\xi(z) \right] \\
 &= \delta v_\xi + (1 - \delta) \left[\int_A u_S(\cdot, \theta) d\rho_\xi(\cdot|\hat{y}_\xi(z), z) - \hat{y}_\xi(z) \right].
 \end{aligned}$$

Hence, her incentive to deviate from message $z \in Z$ to message $z' \in Z$ at such a history in the repeated game is proportional to her incentive to deviate from $(\hat{y}_\xi(z), z)$ to $(\hat{y}_\xi(z'), z')$ in the static game with monetary incentives; and so the S incentive constraints from the latter imply her incentive constraints for the former.

Finally, letting $\tilde{v}_\xi$ denote the S payoff generated by the given strategy profile when starting in state $\xi \in \{0, 1\}$, let us compute $(\tilde{v}_0, \tilde{v}_1)$ and verify that it coincides with (v_0, v_1) . By construction, each $\xi \in \{0, 1\}$ has

$$\begin{aligned}
 \tilde{v}_\xi &= (1 - \delta) \int_{\Theta} \int_Z \int_A u_S(\cdot, \theta) d\hat{\rho}_\xi(\cdot|z) d\hat{\sigma}_\xi(z|\cdot) d\mu_0 \\
 &\quad + \delta \left[\tilde{v}_\xi + (\tilde{v}_{1-\xi} - \tilde{v}_\xi) \int_{\Theta} \int_Z \frac{1-\delta}{\delta} \frac{|\hat{y}_\xi(z)|}{v_1 - v_0} d\hat{\sigma}_\xi(z|\cdot) d\mu_0 \right] \\
 &= (1 - \delta) \int_{\Theta} \int_Z \int_A u_S(\cdot, \theta) d\hat{\rho}_\xi(\cdot|z) d\hat{\sigma}_\xi(z|\cdot) d\mu_0 \\
 &\quad + \delta \left[\tilde{v}_\xi - (\tilde{v}_1 - \tilde{v}_0) \int_{\Theta} \int_Z \frac{1-\delta}{\delta} \frac{\hat{y}_\xi(z)}{v_1 - v_0} d\hat{\sigma}_\xi(z|\cdot) d\mu_0 \right] \\
 &= \delta \tilde{v}_\xi + (1 - \delta) \int_{\Theta} \int_Z \left[\int_A u_S(\cdot, \theta) d\hat{\rho}_1(\cdot|z) - \frac{\tilde{v}_1 - \tilde{v}_0}{v_1 - v_0} \hat{y}_\xi(z) \right] d\hat{\sigma}_\xi(z|\cdot) d\mu_0,
 \end{aligned}$$

which rearranges to

$$\begin{aligned}
\tilde{v}_\xi &= \int_{\Theta} \int_Z \left[\int_A u_S(\cdot, \theta) d\hat{\rho}_\xi(\cdot|z) - \frac{\tilde{v}_1 - \tilde{v}_0}{v_1 - v_0} \hat{y}_\xi(z) \right] d\hat{\sigma}_\xi(z|\cdot) d\mu_0 \\
&= \int_{\Theta} \int_Z \left[\int_A u_S(\cdot, \theta) d\hat{\rho}_\xi(\cdot|z) - \hat{y}_\xi(z) \right] d\hat{\sigma}_\xi(z|\cdot) d\mu_0 \\
&\quad + \frac{(v_1 - v_0) - (\tilde{v}_1 - \tilde{v}_0)}{v_1 - v_0} \int_{\Theta} \int_Z \hat{y}_\xi(z) d\hat{\sigma}_\xi(z|\cdot) d\mu_0 \\
&= \int_{\Theta} \int_{Y \times Z} \left[\int_A u_S(\cdot, \theta) d\rho_\xi(\cdot|y, z) - y \right] d\sigma_\xi(y, z|\cdot) d\mu_0 \\
&\quad + \frac{(v_1 - v_0) - (\tilde{v}_1 - \tilde{v}_0)}{v_1 - v_0} \int_{\Theta} \int_Z \hat{y}_\xi(z) d\hat{\sigma}_\xi(z|\cdot) d\mu_0 \\
&= v_\xi + [(\tilde{v}_0 - v_0) - (\tilde{v}_1 - v_1)] \gamma_\xi,
\end{aligned}$$

where $\gamma_\xi := \frac{1}{v_1 - v_0} \int_{\Theta} \int_Z \hat{y}_\xi(z) d\hat{\sigma}_\xi(z|\cdot) d\mu_0$. Now, because $\gamma_0 \leq 0 \leq \gamma_1$, the pair of linear equations

$$\begin{aligned}
\tilde{v}_1 - v_1 &= \frac{\gamma_1}{(1 + \gamma_1)} (\tilde{v}_0 - v_0) \\
\tilde{v}_0 - v_0 &= \frac{-\gamma_0}{(1 - \gamma_0)} (\tilde{v}_1 - v_1)
\end{aligned}$$

has a unique solution, implying $(\tilde{v}_0, \tilde{v}_1) = (v_0, v_1)$ as desired. ■

We can now easily prove the main theorem. Just as in APS, having established the appropriate monotonicity and closure properties of the relevant set operator, the remaining argument is nearly immediate.

Proof of Theorem 1. We know $V_\delta \neq \emptyset$ because there exists a repeated babbling equilibrium. Given public randomization, V_δ is convex. Let $\bar{v} := \sup V$, $\underline{v} := \inf V$, and $M := \frac{\delta}{1-\delta}(\bar{v} - \underline{v})$. Lemma 6 implies that $V \subseteq G(M)$. Then, Lemma 3 implies that $\bar{V} \subseteq G(M)$. Lemma 7 then implies that $\bar{V} \subseteq V$ (i.e. V is closed), so $V = [\underline{v}, \bar{v}]$. Therefore, V is a compact interval with $V \subseteq G(\frac{\delta}{1-\delta}(\bar{v} - \underline{v}))$. By Lemma 7, any other such interval with the same property is a subset of V . ■

A.3. Equilibrium computations for project implementation

Here, we provide supporting computations for the equilibrium described at the end of Section 4. Our aim is to show that the described strategy profile, paired with some belief map, forms an equilibrium and yields the manager a payoff of $\bar{v}$. Take the CEO's beliefs to be derived from Bayesian updating whenever message z_H or z_L is sent, and have the CEO's belief from any other message be the same as that from z_L ; the Bayesian property is then immediate. The CEO's best response property is also immediate, because message z_H leaves her indifferent between all of her actions, and other messages leave her more pessimistic about the state (and hence willing to choose action zero). Moreover, the manager is always indifferent between sending message z_L

and sending messages outside of $\{z_L, z_H\}$. Hence, all that remains to show is that the manager's payoff, v , from this strategy profile is equal to $\bar{v}$; and that the manager prefers sending z_H to z_L [resp. z_L to z_H] whenever $\theta > \theta_*$ [resp. $\theta < \theta_*$].

Now, let v denote the manager's payoff at the beginning of the game, and let $p := \mathbb{P}\{\theta \geq \theta_*\}$. Note that, given the form of the CEO's strategy, the manager's continuation payoff is $\varepsilon^k v$ at any history for which he has sent message z_H exactly k times. Observe that

$$\begin{aligned} 1 - \varepsilon &= \frac{(1-\delta)\bar{\delta}}{\delta(1-\delta)} = \frac{1-\delta}{\delta} \frac{\theta_*}{\bar{v}} \\ \implies v &= (1-\delta)\mathbb{E}[\theta \mathbf{1}_{\theta \geq \theta_*}] + \delta[1 - (1-\varepsilon)p]v \\ &= (1-\delta)(\bar{v} + p\theta_*) + \delta\left[1 - \frac{1-\delta}{\delta} \frac{\theta_*}{\bar{v}} p\right]v \\ &= (1-\delta)\left[\bar{v} + p\theta_*\left(1 - \frac{\bar{v}}{\theta_*}\right)\right] + \delta v \\ \implies v &= \bar{v} + p\theta_*\left(1 - \frac{\bar{v}}{\theta_*}\right) \end{aligned}$$

The last equation is satisfied by $v = \bar{v}$, and the two sides move antimonotonically with v ; it follows that $v = \bar{v}$. Finally, we verify the manager's incentives. At any history, some $k \in \mathbb{N}$ exists so that the manager's incremental payoff from reporting z_H rather than reporting z_L is

$$\begin{aligned} [(1-\delta)\varepsilon^{k-1}\theta + \delta\varepsilon^k\bar{v}] - \delta\varepsilon^{k-1}\bar{v} &= \varepsilon^{k-1}[(1-\delta)\theta - \delta(1-\varepsilon)\bar{v}] \\ &= \varepsilon^{k-1}\left[(1-\delta)\theta - \delta\left(\frac{1-\delta}{\delta} \frac{\theta_*}{\bar{v}}\right)\bar{v}\right] \\ &= \varepsilon^{k-1}(1-\delta)(\theta - \theta_*). \end{aligned}$$

Hence, types above [resp. below] θ_* prefer to send message z_H [resp. z_L], as required.

A.4. Bonuses can hurt

Our main theorem yields an equivalence between the sender payoffs in the repeated game and the payoffs from a static cheap talk game with money burning and bonuses. However, bonuses played no role in our one worked application. Moreover, it is straightforward to show (for instance, applying an observation of Kamenica and Gentzkow (2011)) that the leading example of Crawford and Sobel (1982) shares the feature that a repeated babbling equilibrium is worst for S in the repeated game, so that again the one-shot game with bonuses is irrelevant. Here, we provide an example in which capped self-assigned bonuses can lead to S achieving worse payoffs in the static game than he could without bonuses.²³

Let $\Theta = \{-1, 0, 1\}$, $A = [-1, 1]$, $\mu_0(1) = \mu_0(-1) = \frac{1}{4}$, and the players' utility functions take the form $u_R(a, \theta) = -(a - \theta)^2$ and $u_S(a, \theta) = a - 12(1 - \theta^2)a^2$. We argue, proceeding informally, that this example is as desired.

First, consider cheap talk with no monetary signals. Facing quadratic loss, R always chooses his expectation of the state. As the two nonzero S types simply want to maximize R's action, there is no equilibrium in which these two types send different message distributions: type -1 would have a strict incentive to deviate from the message she is supposedly more likely to send. Therefore, R chooses action 0 almost surely, yielding S payoff of 0.

Now, we describe an equilibrium of the game where S can send bonuses capped at 1. Here, S collects no bonus if her type is 1, collects a bonus of 1 if her type is -1 , and mixes uniformly

²³ We are grateful to Doron Ravid for providing the first known example to us; this example is essentially the same as Doron's.

between these two choices if her type is 0; she babbles, so that only her bonus signals her type. Meanwhile, R will have uniform belief over types $\{0, 1\}$ and choose $\frac{1}{2}$ if no bonus is collected, and have uniform belief over types $\{-1, 0\}$ and choose $-\frac{1}{2}$ if a strictly positive bonus is collected. Direct computation shows that the above describes an equilibrium (in particular, all S types are indifferent between collecting no bonus and collecting a maximum bonus), and that S has an ex-ante payoff of $-1 < 0$.

References

- Abreu, D., Pearce, D., Stacchetti, E., 1990. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 1041–1063.
- Aliprantis, C.D., Border, K.C., 2006. *Infinite Dimensional Analysis*, second ed. Springer, Berlin.
- Alonso, R., Matouschek, N., 2008. Optimal delegation. *Rev. Econ. Stud.* 75 (1), 259–293.
- Aumann, R.J., Hart, S., 2003. Long cheap talk. *Econometrica* 71 (6), 1619–1660.
- Aumann, R.J., Maschler, M., 1966. Game theoretic aspects of gradual disarmament. *Annu. Rep. - U.S. Arms Control Disarm. Agency* 80, 1–55.
- Aumann, R.J., Maschler, M., Stearns, R.E., 1995. *Repeated Games with Incomplete Information*. MIT Press.
- Austen-Smith, D., Banks, J.S., 2000. Cheap talk and burned money. *J. Econ. Theory* 91 (1), 1–16.
- Best, J., Quigley, D., 2020. Persuasion for the Long Run.
- Chakraborty, A., Harbaugh, R., 2010. Persuasion by cheap talk. *Am. Econ. Rev.* 100 (5), 2361–2382.
- Che, Y.-K., Dessein, W., Kartik, N., 2013. Pandering to persuade. *Am. Econ. Rev.* 103 (1), 47–79.
- Compte, O., 1998. Communication in repeated games with imperfect private monitoring. *Econometrica*, 597–626.
- Crawford, V.P., Sobel, J., 1982. Strategic information transmission. *Econometrica* 50 (6), 1431–1451.
- Dessein, W., 2002. Authority and communication in organizations. *Rev. Econ. Stud.* 69 (4), 811–838.
- Forges, F., 1986. An approach to communication equilibria. *Econometrica*, 1375–1385.
- Forges, F., 1990. Equilibria with communication in a job market example. *Q. J. Econ.* 105 (2), 375–398.
- Fudenberg, D., Kreps, D.M., Maskin, E.S., 1990. Repeated games with long-run and short-run players. *Rev. Econ. Stud.* 57 (4), 555–573.
- Fudenberg, D., Levine, D., Maskin, E., 1994. The folk theorem with imperfect public information. *Econometrica* 62 (5), 997–1039.
- Fudenberg, D., Levine, D.K., 1994. Efficiency and observability with long-run and short-run players. *J. Econ. Theory* 62 (1), 103–135.
- Gilligan, T.W., Krehbiel, K., 1989. Asymmetric information and legislative rules with a heterogeneous committee. *Am. J. Polit. Sci.*, 459–490.
- Glazer, J., Rubinstein, A., 2006. A study in the pragmatics of persuasion: a game theoretical approach. *Theor. Econ.* 1, 395–410.
- Golosov, M., Skreta, V., Tsyvinski, A., Wilson, A., 2014. Dynamic strategic information transmission. *J. Econ. Theory* 151, 304–341.
- Green, J.R., Stokey, N.L., 2007. A two-person game of information transmission. *J. Econ. Theory* 135 (1), 90–104.
- Hart, S., Kremer, I., Perry, M., 2017. Evidence games: truth and commitment. *Am. Econ. Rev.* 107 (3), 690–713.
- Holmström, B., 1982. On the Theory of Delegation. Northwestern University.
- Kamenica, E., Gentzkow, M., 2011. Bayesian persuasion. *Am. Econ. Rev.* 101 (6), 2590–2615.
- Kandori, M., Matsushima, H., 1998. Private observation, communication and collusion. *Econometrica* 66 (3), 627.
- Karamychev, V., Visser, B., 2017. Optimal signaling with cheap talk and money burning. *Int. J. Game Theory* 46 (3), 813–850.
- Kartik, N., 2007. A note on cheap talk and burned money. *J. Econ. Theory* 136 (1), 749–758.
- Kartik, N., Van Weelden, R., 2018. Informative cheap talk in elections. *Rev. Econ. Stud.* 86 (2), 755–784.
- Kolotilin, A., Li, H., 2021. Relational communication. *Theor. Econ.* 16 (4), 1391–1430.
- Krishna, V., Morgan, J., 2004. The art of conversation: eliciting information from experts through multi-stage communication. *J. Econ. Theory* 117 (2), 147–179.
- Krishna, V., Morgan, J., 2008. Contracting for information under imperfect commitment. *Rand J. Econ.* 39 (4), 905–925.
- Levin, J., 2003. Relational incentive contracts. *Am. Econ. Rev.* 93 (3), 835–857.
- Lipnowski, E., Ramos, J., 2020. Repeated delegation. *J. Econ. Theory* 105040.
- Lipnowski, E., Ravid, D., 2020. Cheap talk with transparent motives. *Econometrica* 88 (4), 1631–1660.

- Mailath, G., Samuelson, L., 2006. *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press, New-York.
- Margaria, C., Smolin, A., 2018. Dynamic communication with biased senders. *Games Econ. Behav.* 110, 330–339.
- Mathevet, L., Pearce, D., Stacchetti, E., 2019. Reputation and Information Design.
- Myerson, R.B., 1981. Optimal auction design. *Math. Oper. Res.* 6 (1), 58–73.
- Phelps, R.R., 2001. *Lectures on Choquet's Theorem*, second ed. Springer, Berlin.
- Rappoport, D., 2021. Evidence and Skepticism in Verifiable Disclosure Games.
- Renault, J., Solan, E., Vieille, N., 2013. Dynamic sender–receiver games. *J. Econ. Theory* 148 (2), 502–534.
- Salamanca, A., 2021. The value of mediated communication. *J. Econ. Theory* 192, 105191.
- Spear, S.E., Srivastava, S., 1987. On repeated moral hazard with discounting. *Rev. Econ. Stud.* 54 (4), 599–617.