

Information Transmission under the Shadow of the Future: An Experiment[†]

By ALISTAIR J. WILSON AND EMANUEL VESPA*

We experimentally examine how information transmission functions in an ongoing relationship. Where the one-shot cheap-talk literature documents substantial overcommunication and preferences for honesty, the outcomes in our repeated setting are more consistent with uninformative babbling outcomes. This is particularly surprising, as honest revelation is supportable as an equilibrium outcome in our repeated setting. We show that inefficient outcomes are driven by a coordination failure on how to distribute the gains from information sharing. However, when agents can coordinate on the payment of an “information rent,” honest revelation emerges. (JEL C92, D83)

Economic agents often have the opportunity to seek advice from experts repeatedly (for example: financial advisers, lawyers, mechanics). When interests are not perfectly aligned, economic theory has shown that information transmission can be severely constrained in one-shot interactions (Crawford and Sobel 1982). But long-run relationships open up the possibility for greater efficiency. Our main questions focus on the extent to which repeated interaction produces increased exchange of information, where our findings indicate only a qualified positive, as information is shared only when honesty is compensated. Moreover, we show that solving a coordination problem over how to share efficiency gains is critical to successful outcomes.

Our paper examines the intersection of two key areas of current economic research: the degree to which information is shared by interested parties and whether repeated interactions are leveraged for increased efficiency. Taken independently, the two literatures provide an optimistic prior. First, the experimental literature on one-shot information transmission documents overtransmission relative to the theoretical prediction (with a metastudy in Abeler, Nosenzo, and Raymond 2016).

*Wilson: Department of Economics, University of Pittsburgh, Wesley W. Posvar Hall, Pittsburgh, PA 15224 (email: alistair@pitt.edu); Vespa: Department of Economics, University of California, Santa Barbara, North Hall, Santa Barbara, CA 93106 (email: vespa@ucsb.edu). Johannes Hörner was coeditor for this article. We would like to thank Ted Bergstrom, Andreas Blume, Gabriele Camera, Katie Coffman, Lucas Coffman, John Duffy, Rod Garratt, PJ Healy, Navin Kartik, Andrew Kloosterman, John Kagel, Melis Kartal, Tom Palfrey, Wooyoung Lim, Ryan Oprea, Joel Sobel, Lise Vesterlund, and Sevgi Yuksel. This paper reports on informed consent experiments that were approved by the University of Pittsburgh IRB (PRO-15080059). The authors gratefully acknowledge financial support from the NSF (SES-1629193). The authors work on a number of projects together and have opted to alternate between orders across these projects; author order conveys no information on contribution, which was entirely equal.

[†]Go to <https://doi.org/10.1257/mic.20170403> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

Second, a literature examining repeated interaction in prisoners' dilemmas (see the metastudy in Dal Bó and Fréchette 2018) shows that efficiency is improved through the use of dynamic strategies in repeated prisoners' dilemmas (RPD). However, our examination of repeated cheap talk suggests caution is warranted as we extend to richer settings. With respect to the first finding, informed parties may be discouraged from honestly revealing if decision-makers use the provided information solely to their own advantage. This filters into a problem in extending the second literature: our coordination problem is *not* a trivial extension of the RPD. In an RPD there is typically a unique efficient outcome (joint cooperation), where repeated cheap talk generically produces multiple efficient equilibria dependent on how the informed and uninformed participants divide the informational gains.

In our experiments we use a simple sender-receiver environment with uncertainty over a payoff-relevant state. In the stage game, an investment opportunity is either *good* or *bad* with equal chance. An informed expert (the sender, she) observes the investment quality and provides a recommendation to the decision-maker (the DM, the receiver, he), who chooses an investment level. We restrict the sender to one of two possible recommendations—*invest* or *don't invest*—where her payoffs are increasing in the decision-maker's investment level regardless of the underlying quality. The DM's ex post payoff is maximized by *full investment* in the good state and *no investment* in the bad state. However, in addition to full and no investment, the DM can also select an intermediate *partial* investment level, which is the best response to his prior.

These incentives create a clear strategic tension: completely aligned interests in revealing the good state to produce full investment, but completely opposed preferences if the bad state is revealed. In one-shot interactions the wedge between honest revelation and self-interest for low-quality investments leads to inefficient “babbling” as the unique equilibrium outcome in the stage-game, with the DM choosing the uninformed best-response of partial investment at both states. Repetition allows for efficient equilibrium outcomes via history dependence, similar to the RPD. However, unlike the RPD, a tradeoff emerges over how the surplus generated is distributed. While it is theoretically possible for receivers to achieve their first-best outcome, doing so requires the use of harsh, relatively complex punishments. In contrast, while full revelation is still possible with simple punishments—reverting to the babbling outcome on a deviation—distributionally, it requires an informed receiver to extract less than the myopic best response. The difference in payoffs can be characterized as an “information rent,” a transfer to the sender for honestly revealing. Our paper's findings suggest that miscoordination over this distributional tradeoff severely hampers subjects' ability to achieve efficient outcomes.

Our first experimental finding demonstrates the failure to produce efficiency gains with repeated play. Specifically, a *Partners* treatment with fixed supergame matchings is compared to a *Strangers* control that essentially repeats the anonymous one-shot game. The large majority of subjects in both treatments coordinate on uninformative babbling outcomes. While there are some statistically significant differences driven by repetition, the quantitative effects are small and not economically meaningful. Across a series of experimental manipulations, we then set

out to show that the failure is driven by miscoordination over the payment of an information rent.

Our first modification eases the coordination problem by providing subjects with preplay communication, free-form *Chat* before the supergame begins. *Chat* increases the efficiency by 58 percentage points relative to Partners. Moreover, the exchanged messages clearly indicate that the large majority of subjects are exactly coordinated on an information rent payment. In a second Partners modification (without preplay communication), we allow receivers to pay an information rent through an explicit side payment rather than through a myopically suboptimal action. While not as successful as preplay communication, making the rent payment explicit increases efficiency gains by 24 percentage points relative to Partners. Finally, in two further robustness treatments, we modify the Partners' game payoffs. While both changes make the games less economically representative for our sender-receiver question, the changes make the payment of an information rent unnecessary, and full extraction by receivers is possible in equilibrium with a simple babbling punishment. The results in both robustness treatments indicate that removing the requirement to pay an information rent removes the coordination hurdle subjects stumble upon in Partners.

Stepping back from the sender-receiver game, our paper documents a failure to solve a novel coordination problem over dynamic strategies—distinct from the typical coordination problem in the workhorse RPD. In our setting, players must also coordinate over which efficient outcome is selected, a distributive tension over how the gains are split. Moreover, solving this problem is directly related to the severity of the dynamic strategy punishments used to support any division. While harsh, minmax-like punishments can support one party getting the lion's share, we do not find any evidence for these strategies. In contrast, restriction to simpler Nash reversion punishments not only provides a better match to punishment behaviors, the restriction also predicts the on-path information rents paid in our successful partnerships.

Restricting ourselves to repeated cheap talk, the evidence across our experimental treatments points to the importance of information rent payments, an idea that goes back to an older literature on regulating a monopolist (see Baron and Besanko 1984, Laffont and Tirole 1988). Overall, our findings mirror the idea that honest advice must be paid for in the long run and must be commonly known to be paid for. Decision-makers that fail to explicitly reward those providing honest information today degrade the extent to which they get honest advice tomorrow.

I. Literature

The paper directly relates to two established literatures: information transmission and repeated games. For the first, our stage game is related to cheap talk games examining an informed but biased sender, where the most prominent paper in this literature is Crawford and Sobel (1982). The fundamental theoretical finding is that full revelation is ruled out (for the one-shot game) with at best partial revelation in equilibrium. The Crawford-Sobel model has been examined in a number of experimental studies (Dickhaut, McCabe, and Mukherji 1995; Cai and Wang

2006; Wang, Spezio, and Camerer 2010) where the main results are that more information is exchanged than predicted by theory.¹ Outside of the Crawford-Sobel framework, a number of papers provide evidence consistent with honest revelation in cases in which lying is profitable (for example, Gneezy 2005, Fischbacher and Föllmi-Heusi 2013).² Our findings suggest that in a repeated setting, a necessary (though not sufficient) condition for honest revelation is that honesty is rewarded.³ This is similar to the difference between unconditional cooperation documented for one-shot PD and conditional cooperation found for the RPD (Dal Bó and Fréchet 2011), the workhorse stage game in theoretical and applied work. However, there are two qualitative differences between the RPD and our repeated cheap talk stage game.

First, while there is a unique efficient outcome in the typical RPD parameterization, there are multiple efficient outcomes in the environment we study. Agents who want to reach an efficient outcome need to solve a distributional problem that is absent in an RPD. From the perspective of distributional difficulties, the experimental literature that is closest to our setting is the repeated trust game.⁴ In each period, a trustor decides between *out*, where both trustor and trustee receive inefficient outcomes, and *in*, which is efficient but triggers a distributional choice from the trustee. The experimental literature in the area suggests that agents can leverage the infinite time horizon to cooperate, where Engle-Warnick and Slonim (2004, 2006) find that efficient outcomes are reached more often with an indefinite horizon (relative to fixed horizon), with the potential for highly efficient and equitable outcomes. In contrast, our information transmission game requires trust in both directions. A revealing bad-state message has the sender “trust” the receiver to make a distribution. Meanwhile, facing a message signaling the good state introduces the issue of trust in the opposite direction: where the receiver needs to decide whether to “trust” the sender or not. Our findings in the Partners treatment suggest that the results from repeated trust games do not extend to our setting in which distributional issues take place under incomplete information.

A second difference with the RPD is the available punishments. In the RPD the worst punishment that can be used to leverage cooperation (the *minmax*) is identical to the stage game Nash. In contrast, receivers in our game have access to harsher punishments than Nash reversion. In an RPD meta-analysis, Dal Bó and Fréchet (2018) suggest that most subjects use punishments that are at the stage game Nash

¹ See also Blume et al. (1998) for an examination of the evolution of message meaning, and Lai, Lim, and Wang (2015) and Vespa and Wilson (2016) for the extension to multiple senders.

² Some of the excess revelation in Gneezy can be explained as failed strategizing by senders believing the receiver will reverse their advice (see Sutter 2009), though only 20 percent of receivers do so. See also Hurkens and Kartik (2009), which demonstrates that lying aversion may not vary with the harm done to others. Closer to our findings, Le Quement and Patel (2018) discusses the constraints on honesty in a static setting with reciprocity.

³ In our repeated setting, the state is randomly drawn each period, which differs from recent theoretical work examining dynamic sender-receiver environments with persistent state variables (see Renault, Solan, and Vieille 2013 and Golosov et al. 2014). Experimentally, Ettinger and Jehiel (2016) examines a repeated sender-receiver game where the sender has a persistent type, mirroring the Sobel (1985) model of credibility. See also the repeated sender-receiver game of Kartal, Müller, and Tremewan (2017), which experimentally studies the effects of gradualism on building trust. For a connection to information design, see Mathevet, Pearce, and Stacchetti (2019).

⁴ More broadly, there is a large literature studying bargaining situations in which efficiency considerations may be absent. See Roth (1995) for a survey of experiments on bargaining, Goeree and Holt (2000) for an experimental study of alternating bargaining offers where (as in our setting) payoffs across players are not symmetric, Güth and Kocher (2014) for a recent work on the ultimatum game, and Fréchet and Vespa (2017) for multilateral bargaining.

(via grim trigger strategies) or are less harsh than the stage game Nash (with tit for tat a frequently detected strategy).^{5,6} However, because the minmax and stage game Nash outcomes coincide, the evidence from the RPD cannot assess whether harsher-than-Nash punishments have positive content. From a theoretical perspective, minmax punishments can be used to leverage more cooperative outcomes (Fudenberg and Maskin 1986). While such strategies constitute equilibria of our game, we find that out of all available punishments, only simple punishments (in particular, Nash reversion) are used. Moreover, the additional predictive power that comes from focusing on outcomes supported with simple punishments leads theoretically to the main finding in our experiments: the on-path payment of an information rent is required to support efficient outcomes. To our knowledge, our paper is the first to examine the tradeoffs between sharing the gains from the selection of more efficient equilibria and the selection of punishments.⁷

II. Design: Theory and Hypotheses

In this section we present the laboratory environment, a streamlined version of a repeated cheap talk game. The stage game of the sender-receiver environment we study is depicted in Figure 1, panel A. The game unfolds as follows: (i) Nature chooses a state of the world $\theta \in \{\text{Good}, \text{Bad}\}$, with equal probability; (ii) the *sender* perfectly observes the state θ and chooses a message $m \in \{\text{Invest}, \text{Don't Invest}\}$; (iii) the *receiver* observes the message m but not the state and chooses an action $a \in \{\text{Full}, \text{Partial}, \text{None}\}$; (iv) at the end of each stage game, the entire history of the round (θ, m, a) becomes common knowledge.⁸

The sender's payoff is

$$u(a) = \begin{cases} 1 & \text{if } a = \text{Full} \\ 1/3 & \text{if } a = \text{Partial} \\ 0 & \text{if } a = \text{None}, \end{cases}$$

⁵In a similar vein, punishment phases in imperfect public monitoring games are typically more forgiving and less harsh than the static Nash outcome (Fudenberg, Rand, and Dreber 2012). In dynamic game settings with state variables, the most frequent punishments are reversions to the Markov perfect equilibrium (Vespa and Wilson 2017, 2019). More costly punishments than Nash reversion have been studied, though none find that this selection is useful in supporting higher payoffs. Dreber et al. (2008) finds that where costly punishments are used, they do not succeed at supporting better outcomes. Wilson and Wu (2017) finds subjects employing a dominated strategy (termination) as a costly punishment but without an observed increase in cooperation.

⁶For the effect of preplay communication in repeated PD games, see Arechar et al. (2017) and references therein. Also see the literature on team communication, in particular Cooper and Kagel (2005) and Kagel and McGee (2016).

⁷Outside of the infinite horizon setting, Brown, Falk, and Fehr (2004) shows that gains over the inefficient equilibrium prediction need to be shared for efficient outcomes to emerge in a relational contracting setting. However, if placed in an infinite horizon setting, these results can be rationalized with simple punishments based on termination and rematching.

⁸In the laboratory we use a neutral frame where the state, messages, and recommendations are labeled as either "left" (for good/invest/full investment) or "right" (for bad/don't/no investment), with "middle" the label for the safe partial-investment action. Instructions are available in online Appendix G.




FIGURE 1. STAGE GAME

while the receiver’s payoff is

$$v(a, \theta) = \begin{cases} 1 & \text{if } (a, \theta) \in \{(Full, Good), (None, Bad)\} \\ 2/3 & \text{if } a = Partial \\ 0 & \text{otherwise.} \end{cases}$$

Conditional on a good state, the two players share common interests: the investment actions are Pareto ranked (from best to worst) as Full, Partial, None. However, conditional on a bad state, the sender and receiver have diametrically opposed interests. In particular, when $\theta = Bad$, a Full investment choice distributes everything to the sender, while None distributes everything to the receiver. Finally, we choose an asymmetric distribution on Partial investment so that it is preferred by an uninformed receiver to either full or no investment.⁹

The informed sender’s chosen signal m affects neither players’ payoff directly, so this is a cheap-talk environment. As the sender’s payoff is state independent (affected by the chosen action only), the theoretical prediction for the stage game is that every perfect Bayesian equilibrium (PBE) involves the receiver choosing $a = Partial$ with certainty.¹⁰

A. Experimental Design

Our experimental design examines a repeated (infinite horizon) version of the above stage game. After each completed round, we continue the supergame

⁹This is under risk neutrality, where a posterior belief on either state above 2/3 is required for partial investment not to be the best response. For partial not to be the best response if uninformed, the receiver must be risk loving to a very strong degree (a CRRA coefficient of -0.71 or below).

¹⁰The outcome in which the receiver takes the uninformed best response is referred to as babbling. One reason we introduce a third action for the receiver is to make the uninformed best response unique.

into another round with probability $\delta = 3/4$, and with probability $(1 - \delta)$, the supergame ends. From the point of view of a risk-neutral subject in round one, the expected payoffs for the sequence of possible states/messages/actions for the sender and receiver, respectively, is proportional to

$$U(\{m^t(\theta), a^t(m)\}_{t=1}^{\infty}) = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} E_{\theta} u(a^t(m^t(\theta))),$$

$$V(\{m^t(\theta), a^t(m)\}_{t=1}^{\infty}) = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} E_{\theta} v(a^t(m^t(\theta)), \theta^t).$$

We initially examine two treatments. In our *Partners* treatment, the sender-receiver pair is fixed across the supergame. A sender participant i and a receiver participant j are paired together in all supergame rounds, with all uncertainty realized at the end of each round's stage game. Senders in round $t > 1$ observe their matched receiver's actions choice in each prior rounds; receivers observe the actual state realization.

In our second treatment, which we call *Strangers*, the sender-receiver pairings in each round are random. A particular sender i is anonymously matched with a random sequence of receivers, $J(t)$, across the payment supergame. Similarly, each receiver j is matched to a random sequence of senders, $I(t)$. While observing their own supergame history, each sender (receiver) has no knowledge of the prior play of the receiver (sender) they are matched to in each new round.¹¹

Strangers Game.—Given the random anonymous matching, we treat each *Strangers* treatment supergame as a sequence of one-shot games. Though the instructions, supergame lengths, feedback, and potential payoffs are the same as our *Partners* treatment, and subjects are unable to tie the previous history to the presently matched participant. The shadow of the future cannot work if subsequent behavior is independent. In principle, were δ large enough, more efficient equilibria are possible in *Strangers* through “contagion” constructions. We do not focus on such constructions for two reasons: First, experimental examination of contagion equilibria shows that they lack predictive power when they do exist.¹² Second, given our choice of $\delta = 3/4$ and the random matching protocol we use, contagion constructions do not allow for efficient, fully revealing equilibria in the *Strangers* treatment.¹³

Given that the *Strangers* game is effectively just many repetitions of the one-shot game, the theoretical prediction for *Strangers* treatment is simply the stage game prediction: babbling in all rounds.

¹¹ Moreover, we guarantee senders and receivers in *Strangers* that the rematching process is chosen to make sure that they are never rematched to the same subject in two contiguous rounds: so $\Pr\{I(t) = I(t+1)\} = \Pr\{J(t) = J(t+1)\} = 0$ for all t .

¹² Duffy and Ochs (2009) finds no support for contagion equilibrium predictions. Camera and Casari (2009) shows that contagion can be observed if subjects are provided the history of their anonymous match's choices and the group size is small (the strongest evidence is for groups of four). Neither of these conditions are met in our design, where subjects do not observe their match's previous choices and the group size in our sessions consists of at least 14 participants.

¹³ See online Appendix A for a proof.

RESULT 1 (Strangers): *All equilibrium outcomes have the following features: (i) the sender's message does not reveal the state, and (ii) the receiver selects the inefficient uninformed best response (Partial).*

A PBE for the stage game requires us to specify a sequentially rational behavioral strategy $\beta(a|\mathcal{H})$, a probability of selecting each available action a at each information set \mathcal{H} , and a system of beliefs $\mu(\theta|\mathcal{H})$ that is consistent at every information set reached with positive probability. A full specification of an intuitive babbling PBE has a sender choice of $\beta(\text{Invest}|\theta) = 1$ for both states, a receiver choice of $\beta(\text{Partial}|\text{Invest}) = 1$ and $\beta(\text{None}|\text{Don't}) = 1$, with $\mu(\text{Good}|\text{Invest}) = 1/2$ and $\mu(\text{Good}|\text{Don't Invest}) = 0$.¹⁴

Partners Game.—With fixed supergame partners, more informative outcomes than the stage game PBE are supportable as equilibria of the supergame so long as δ is large enough. Figure 1, panel B shows the set of feasible (both shaded regions) and individually rational (dark gray region) expected payoffs for the stage game. The efficient frontier runs from the extreme point $(1/2, 1)$ labeled *Full Extraction* (as the receiver achieves their first best, matches their action to state) to the extreme point $(1, 1/2)$ labeled *Always Invest* (receiver always chooses Invest, the sender's first best). Note that by construction, a payoff in our game is on the efficient frontier if and only if the receiver fully invests in the good state. Investment behavior in the bad state therefore traces out the points on the frontier, where our constant sum bad-state payoffs lead to the one-for-one downward gradient.

Away from the frontier, the figure illustrates the payoffs from the stage game's $(1/3, 2/3)$ babbling outcome payoff, coinciding with the receiver's individually rational payoff, as they have decision-making power. In contrast, the sender has no decision-making power, and so her individually rational payoff is lower than babbling, as the receiver can enforce a zero payoff by choosing no investment. As such, the worst-case individually rational joint payoff is the point $(1/12, 2/3)$.¹⁵

With fixed pairs, history-dependent strategies can be used to support any individually rational payoffs so long as δ is large enough via standard folk theorem arguments.¹⁶ Restricting our attention to simple pure strategy outcomes, the inefficient, stationary PBE still exists. However, in addition to babbling, there are history-dependent equilibria at $\delta = 3/4$ that support other outcomes on the efficient frontier—in particular, the points *Full Information* and *Information rent* in Figure 1, panel B can both be supported by pure action history-dependent strategies (constructions below).

¹⁴ While there are many different PBE constructions for this game—in particular with different off-path constructions, where some can reveal partial information about the state—all PBE share the feature that Partial investment is chosen in both states.

¹⁵ This payoff is obtainable in the stage game through a fully informed receiver choosing no investment in the bad state and randomizing equally between no and partial investment when the state is good.

¹⁶ Given a public randomization device, all payoffs in the individually rational set are attainable as equilibria at $\delta = 3/4$. This was verified with the Java tool “rgsolve” (Katzwer 2013), which uses the Abreu and Sannikov (2014) algorithm, which we use to examine the stage game in its normal form.

RESULT 2 (Partners): *While the history-independent prediction is still babbling, there are efficient history-dependent outcomes where the sender fully reveals the state.*

The differences across our two treatments can therefore be summarized in the following hypothesis.

HYPOTHESIS 1: *Information transmission and efficiency in the Partners game should be equal to or higher than in the Strangers game, where any information transfer is supported through history-dependent play.*

While repetition has the potential to improve outcomes here, we now show that the coordination hurdles in attaining efficient play in this setting are nontrivial. In particular, supporting full revelation requires either the on-path payment of an “information rent” by the receiver or the use of a harsh, complex punishment.

We first construct the equilibrium in which the sender is paid an information rent. Consider the set of efficient, individually rational payoffs in Figure 1, panel B, the sender-receiver payoff profiles in $\left\{ \frac{1}{2}(1, 1) + \frac{1}{2}(\alpha, 1 - \alpha) \mid 2/3 \geq \alpha \geq 0 \right\}$. Efficiency requires the receiver to choose Full in the good state, while in the bad state, the total unit surplus can be split, with α going to the sender and $(1 - \alpha)$ to the receiver. So long as $\alpha \leq 2/3$, the receiver’s overall payoff will exceed or match his individually rational payoff. However, whenever the bad state is revealed by the sender, the receiver’s myopically rational choice is None, with $\alpha = 0$ to the sender. Because of this, we refer to any bad-state sender payoff with $\alpha > 0$ as an information rent. In equilibrium, any sender payoff of $\alpha > 0$ must be predicated on future revelation.

RESULT 3: *In the Partners game, all fully revealing efficient equilibrium supported by babbling reversion provide the sender with a positive information rent.¹⁷*

The result tells us that sender honesty in our game must be rewarded if the simplest coordination-free punishment is used. While this is true in our particular game, this is a generic property for all $\delta < 1$ in sender-receiver constructions where the sender has a concave state-independent utility function over actions, and a proof of this statement is provided in online Appendix A.

A particularly simple separating equilibrium with information rents is one in which the receiver selects Partial upon receipt of the Don’t Invest message, providing an information rent of $1/3$ to the sender. We will refer to this pure strategy construction as the “information rent” equilibrium.

Result 3 focuses on behavior supported by a simple history-dependent punishment strategy: babbling reversion. However, if receivers are able to coordinate with the sender on the use of harsher punishments, it is possible for them to avoid paying an information rent.

¹⁷ Given full revelation and babbling reversion, a sender deviation in the bad state yields a discounted average payoff of $(1 - \delta) + \delta/3$. If $\alpha = 0$, the on-path payoff is $\delta/2$, so the deviation is profitable whenever $\delta < 6/7$.

RESULT 4: *Receivers can attain their first-best outcome in equilibrium for the Partners game through harsher punishments than babbling.*¹⁸

While repeated interaction allows for the possibility that receivers get their first-best outcome, not paying an information rent comes at a cost in terms of strategic complexity. If senders are to be persuaded to fully reveal without an information rent, they must be threatened with harsher punishment than babbling. Receivers must temporarily choose a myopically suboptimal response (*None*), and for this to be incentive compatible, the receiver must get some compensating benefit from additional revelation in the future, necessitating coordination with the sender. The requisite punishment is therefore not only complex in length, it also requires strategic coordination among the participants. In contrast, a babbling punishment can be unilaterally imposed by either party.

Results 3 and 4 set up our hypothesis on behavior within the Partners game. Should efficient, honest revelation emerge, theory places additional restrictions on how it is supported.

HYPOTHESIS 2 (Supporting Revelation): *If senders reveal information in the Partners game, one of the following holds:*

- (i) *Receivers fully extract but punish dishonesty with no investment for a number of periods before reCOORDINATING on revelation.*
- (ii) *Dishonesty is punished by babbling reversion, but receivers provide an on-path information rents to senders.*

Summary and Session Details.—For both the Strangers and Partners treatments, we recruited a total of 46 unique subjects across 3 separate sessions (16, 16, and 14 subjects).^{19,20} A session consists of 20 repetitions of the supergame, where each repetition involves an unknown number of rounds. (In each round the game continues for one more round with probability $\delta = 3/4$.)²¹ Half of the subjects are randomly assigned to be senders and half receivers, where role assignments are then fixed throughout the session. In the Strangers (Partners) treatment, subjects were randomly matched to a different participant at the beginning of each new round (supergame). Stage game payoffs in the laboratory are multiplied by three so that the bad state involves a constant sum game over \$3, and in the good state, the efficient

¹⁸For proof, consider a construction with a two-phase punishment strategy following any *Sender* deviation from full revelation (or *Receiver* deviation from the punishment path) where the *Receiver* chooses *None* in the next k_1 rounds, then *Partial* in the following k_2 rounds, while the *Sender* babbles for $k_1 + k_2$ rounds. After the $k_1 + k_2$ punishment rounds both, sender and receiver revert to the fully extractive path. Full revelation/extraction supported by this punishment is a PBE at $\delta = 3/4$ for any $(k_1, k_2) \in \{(2, 4), (2, 5), (2, 6), (3, 2), (3, 3)\}$.

¹⁹All sessions were conducted at the University of Pittsburgh, and we used zTree (Fischbacher 2007) to construct our interface. See online Appendix G for screenshots.

²⁰Data, experimental software, and analysis scripts are provided in “Replication Data for: Information Transmission under the Shadow of the Future,” American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. <https://doi.org/10.3886/E115922V1>.

²¹Supergame lengths and the random state draws are fixed by session so that treatment comparisons are balanced over realizations.

TABLE 1—EXPERIMENTAL TREATMENT SUMMARY

	Baseline		Robustness ₁		Robustness ₂	
	Good	Bad	Good	Bad	Good	Bad
<i>Panel A. Payoff tables used</i>						
Full	\$3, \$3	\$3, \$0	\$3, \$3	\$3, \$0	\$3, \$3	\$3, \$0
Partial	\$1, \$2	\$1, \$2	\$1, \$2	\$1, \$2	\$0, \$2	\$0, \$2
None	\$0, \$0	\$0, \$3	\$1, \$0	\$1, \$3	\$0, \$0	\$0, \$3
	Strangers	Partners	Chat	Transfer	Partners-R ₁	Partners-R ₂
<i>Panel B. Differences across treatments</i>						
Stage game payoff table	Baseline	Baseline	Baseline	Baseline	Robustness ₁	Robustness ₂
Supergame matching?	Random	Fixed	Fixed	Fixed	Fixed	Fixed
Preplay communication	None	None	SG 13–20	None	None	None
Sender transfer	None	None	None	\$1 option	None	None
Subjects	46	46	48	46	46	46

outcome allocates \$6 in total.²² While we introduce further treatments in Section IV, Table 1 provides a summary description of all treatments. Panel A shows stage game payoffs, and panel B indicates how treatment differ from one another.

III. Empirical Analysis: Partners and Strangers

In what follows, we present our main findings for Partners and Strangers, which we characterize as a *qualitative null* result. Despite some weakly significant differences in the predicted direction, both treatments are mostly coordinated on the stage game equilibrium (babbling). Table 2 provides a summary of average behavior within our experimental stage games.²³ The table provides empirical analogs to the components of the stage game PBE: (i) $\hat{\beta}(x|\mathcal{H})$, the average rate at which subjects choose response x at information set \mathcal{H} ; (ii) $\hat{\mu}(\theta|m)$, the empirically consistent belief for a receiver, computed as the relative frequency of each state θ conditional on message m ; and (iii) our efficiency measure, the rate at which full investment is chosen in the good state, $\hat{\Pr}(\text{Full}|\text{Good})$. In total, our data contain 460 supergames for each treatment. The table breaks the choices in these supergames out by the initial response (the first round of each supergame, $t = 1$) and overall across all supergame rounds.

The first set of results in Table 2 outline the sender’s strategy, where we provide the proportion of honest responses (using a natural language interpretation) after each state is realized.²⁴ The first row indicates that senders choose the Invest message in the good state at high rates (above 97 percent). This is true in both the

²²Each session took approximately 90 minutes. Payment for the session consisted of a \$6 show-up fee plus the payoff that subjects received in three randomly selected supergames. Including the sessions we introduce in Sections 4 and 5, average payments for subjects were \$28.32, with a minimum of \$7 and a maximum of \$87.

²³A coding mistake led to an error in the subject matching in rounds two onward in supergame 20 of the first Partners session. Even though subjects were unaware of this, we drop the relevant data for this supergame. Tables C4 and C5 in online Appendix C provide results restricted to the last eight supergames.

²⁴The experiment framing for the state is either “L” or “R” where the selected message is either “L” or “R,” so the experimental natural language framing is stronger than the paper’s economic framing.

TABLE 2—BEHAVIORAL RESPONSE, EMPIRICALLY CONSISTENT BELIEFS AND OUTCOMES
(*Strangers* VERSUS *Partners*)

Category	Measure	Strangers		Partners	
		Rd.1	All	Rd.1	All
Sender behavior	$\hat{\beta}(\text{Invest} \text{Good})$	0.992 (0.006)	0.983 (0.005)	0.996 (0.004)	0.976 (0.007)
	$\hat{\beta}(\text{Don't} \text{Bad})$	0.259 (0.029)	0.198 (0.019)	0.422 (0.032)	0.316 (0.021)
Consistent belief	$\hat{\mu}(\text{Good} \text{Invest})$	0.572 (0.010)	0.551 (0.004)	0.633 (0.013)	0.588 (0.008)
	$\hat{\mu}(\text{Bad} \text{Don't})$	0.968 (0.022)	0.922 (0.020)	0.990 (0.010)	0.929 (0.020)
Receiver behavior	$\hat{\beta}(\text{Full} \text{Invest})$	0.360 (0.024)	0.272 (0.017)	0.412 (0.026)	0.310 (0.019)
	$\hat{\beta}(\text{Partial} \text{Invest})$	0.608 (0.024)	0.685 (0.018)	0.569 (0.026)	0.635 (0.020)
	$\hat{\beta}(\text{None} \text{Don't})$	0.567 (0.063)	0.544 (0.039)	0.673 (0.048)	0.600 (0.035)
	$\hat{\beta}(\text{Partial} \text{Don't})$	0.383 (0.062)	0.394 (0.039)	0.276 (0.046)	0.346 (0.036)
Efficiency	$\hat{P}r(\text{Full} \text{Good})$	0.356 (0.031)	0.271 (0.020)	0.430 (0.033)	0.321 (0.023)

Note: Standard errors (in parentheses) are derived from a bootstrap (of size 5,000) with supergame-level resampling, stratified by treatment.

Strangers and Partners treatments, particularly for the initial response where the Invest message is selected in over 99 percent of supergames.

The strategic tension for senders arises when the state is bad, where revelation may lead to a zero payoff. The “honest” response in the bad state is to send the Don’t Invest message; however, our results indicate that this is not the modal selection. In the first round of Strangers supergames, senders select the Don’t message in the bad state 26 percent of the time. This declines to 20 percent when we look at all rounds. For the Partners treatment, the rate of honest revelation of the bad state does increase. However, at 42 percent in the first round of each supergame, honesty is not the most common response, with 58 percent of senders instead opting for the Invest message in the bad state.

Honest revelation is *not* the norm in our experimental sessions. This is true for both the Strangers and Partners treatments, for the behavior at the start of each supergame and overall, at the start of our sessions and at the end.^{25,26}

Though honesty is not the norm in the bad state, the Invest message does convey some information. The next set of results in Table 2 makes precise how much. Each

²⁵ In the first round of the first supergame in our Partners (Strangers) treatment, the message Don’t is sent following a bad state in 33.3 (41.7) percent of our data. In contrast, the message Invest is sent 100 (100) percent of the time following the good state in the very first round of the session.

²⁶ Looking at the last five supergames of each session, the proportion of Don’t messages when the state is bad are 39 and 23 percent for the Partners and Strangers treatments, respectively. If anything, senders reveal information less often as the session proceeds.

entry for $\hat{\mu}(\cdot)$ provides the empirical likelihood for the state conditional on each message, the belief consistent with the data. While the prior is 50 percent on each state, after receiving an Invest message, the good state is more likely in both treatments. At its highest point, the consistent belief on Invest indicates a 63.3 percent probability of the good state in the first round of Partners. The expected payoff to a receiver choosing full investment is therefore \$1.89, just slightly lower than the \$2.00 receivers can guarantee themselves by choosing partial investment.²⁷

In contrast to the Invest message, given a Don't message, the probability of a corresponding bad state is very high (99 percent in the first round of Partners' supergames). The expected round payoff to the receiver choosing no investment after a Don't message is therefore \$2.97 in Partners. The higher payoff from no investment following Don't extends to the Strangers treatment with an expected payoff of \$2.77. Across both treatments, the empirically consistent beliefs indicate that receivers' myopic best response is to choose partial investment following an Invest message and no investment following Don't.

The final two groups of results in Table 2 indicate how receivers in our experiments behave and the ensuing efficiency. Given three actions for each receiver information set (the Invest and Don't messages), we provide the proportion of choices that match the message meaning (full and no investment, respectively) and the proportion of partial investment choices. In particular, full investment in the good state is necessary for an efficient outcome but involves trading off a certain payoff of \$2 against a lottery over \$3 and \$0, with probabilities dependent on the belief the sender is truthfully revealing. Below the receiver's behavioral response to a potentially dishonest message, we provide the simplest proxy in our data for the overall efficiency of play, how often the receivers select full investment in the good state.

The table makes clear that the modal response for receivers in both Strangers and Partners mirrors the myopic best response: partial investment in response to Invest messages and no investment in response to Don't Invest messages. Though the most common response to Invest is incredulity, a large fraction of receivers do choose to respond by investing fully. This is true for 41 percent of decisions in the first round of a Partners supergame and 36 percent of decisions in the first round of Strangers supergames. This difference leads to a marginally significant increase in the overall efficiency for the Partners treatment, where 32.1 percent of good-state rounds have full investment, relative to 27.1 percent in Strangers.

Turning to the distributional outcomes, the majority of receivers respond to the Don't message with no investment; however, a minority do choose partial investment even though the bad state has a more than 90 percent likelihood at this choice node. Remarkably, the data indicate a higher rate of partial investment in the Strangers treatment. Subjects are therefore more likely to make choices corresponding to the payment of an information rent in a one-shot interaction than the repeated relationship.

The evidence from the two treatments points to the following finding.

²⁷For details on discounted average payoffs, see Table C1 in online Appendix C.

FINDING 1: *Aggregate behavior in Strangers mirrors the babbling outcome. While aggregate behavior in our Partners treatment does show greater revelation than Strangers, the difference is quantitatively small, and the response is still best characterized as babbling. Overall, there is only a 5 percentage point increase in efficiency in Strangers relative to Partners.*

While there is the possibility for substantial efficiency gains in Partners through history-dependent strategies, there appears to be only a small treatment effect. In online Appendices D and E, we provide evidence that there *are* significant differences in behavior across these two treatments in the expected direction once we examine the dynamic responses. In particular, we provide evidence for history-dependent play in Partners that is absent in the Strangers treatment. However, though we are able to show statistical differences across Partners and Strangers in the direction of Hypothesis 1, the main economic finding is essentially a null result.

From the behavior reported in the literature on repeated PD games, we would expect the Partners treatment to be highly efficient, with a substantial gap over Strangers. Similarly, the prior literature on one-shot sender-receiver games would lead us to expect excess honesty in both the Strangers and Partners treatments. Instead, efficiency is low in both treatments, with only economically small differences across them relative to the potential gains. Given the failure of Hypothesis 1, the conditional predictions in Hypothesis 2 seem moot. However, in a series of additional treatments, we probe further. In particular, we identify the precise driver for our null result as the failure to coordinate on and interpret the payment of an information rent, necessitated by the babbling reversion punishment—as outlined in Hypothesis 2.

IV. Identifying the Source of Efficiency Failures

Why do subjects fail to achieve better outcomes in the Partners treatment? One explanation is that the representative participant does not understand *how* to formulate a dynamic strategy that supports higher payoffs than babbling. This would be the case if the information rent equilibrium strategy was simply too difficult for subjects to consider. Another option is that subjects do understand the alternative strategies that support efficiency, but they miscoordinate on the chosen efficient equilibrium. This could happen if both parties are trying to coordinate on their own best outcomes—receivers aiming at full extraction, senders for an information rent. Yet a third option is that both sides are approximately coordinated on the same objectives but find the tacit implementation too ambiguous. This would be that case if both sides wanted to implement an information rent outcome but were unsure that the other had the same goal. For example, a receiver observing an initial message indicating the good state may be conflicted, thinking it equally likely that the sender was coordinated on either full revelation or babbling. A wait-and-see approach might lead a receiver to choose partial investment—causing an initially honest sender to revert to babbling. In fact, even when senders indicate their choice by revealing the bad state with a Don't Invest message, a receiver's intention in selecting partial investment as a rent might still seem ambiguous to the sender. Did they choose partial as a reward for honesty or as an unconditional babbling response?

In this section we use a number of control treatments to further identify the reasons behind the low efficiency outcomes in Partners. The manipulations we introduce demonstrate that subjects are indeed aware that it is possible to achieve better outcomes. Not only is the information rent strategy understood by both senders and receivers as a viable option, it is also jointly acceptable to both sides as compromise. Moreover, we show that even a small coordinative push that clearly separates the payment of an information rent from the response to information leads to much greater efficiency. Finally, modifying the game's incentives to remove the requirement for an information rent, the results are again coordinated on efficient play at qualitatively similar levels to those in comparable RPDs.

A. Chat

Our first Partners modification fully relaxes any coordination difficulties, providing the opportunity for the exchange of preplay messages. While identical to the Partners treatment for the first 12, in the last 8 supergames of our 3 *Chat* sessions, the subjects were provided with a free-form chat interface to communicate with their upcoming supergame partner.^{28,29} If subjects are aware of the potential for more efficient play but find it too hard to tacitly coordinate, the provision of a device for explicit strategic coordination will relax this problem, and so we should see efficiency gains over the Partners sessions.

The first two data columns in Table 3 indicate aggregate behavior for the Partners and Chat treatments in the last eight supergames. Though behavior in the first 12 supergames is inseparable across the 2 treatments, stark differences emerge once we provide the coordination device.³⁰ Overall, Chat senders report the truth more than 80 percent of the time, where their honesty in the bad state is increased by almost 50 percentage points over Partners. The effect of greater honesty in the bad state is that messages indicating the good state are more credible. Examining the empirically consistent belief $\hat{\mu}(\theta = \text{Good} | m = \text{Invest})$, an Invest message is 25 percentage points more likely to correspond to the good state in Chat than Partners, enough that the aggregate best response for receivers is now full investment. Looking at the choices of receivers in Chat shows their response to the greater information content of an Invest message, with full investment chosen 83 percent of the time, compared to just 22 percent in Partners.

Providing a strong coordination device therefore leads to large efficiency gains. Conditional on a good-state realization, full investment is selected 28.4 percent of

²⁸ At the beginning of the session, subjects are told that the experiment consists of two parts, where Part I involves 12 supergames. Only when the instructions for Part II are read (prior to supergame 13) are subjects informed that they will be allowed to chat in the last 8 supergames. In the online Appendix, we provide a comparison of Partners and Chat treatments for the first 12 supergames, showing no significant differences.

²⁹ In each Chat session, we recruited 16 subjects (so 48 total) so that we could use a perfect stranger matching protocol for the last 8 supergames. The chat interface is available to the matched sender-receiver pairs for two minutes before the supergame begins, but it is not available once the first stage game begins. As such, chat provides an opportunity for the subjects to coordinate on supergame strategies but does not provide a channel for transmitting messages on state realizations or for strategic recoordination in the evolving supergame.

³⁰ Significant differences in the table are based on *t*-tests with supergame-level observations. In Table D19 in the online Appendix, we provide nonparametric tests, which show significant differences across honesty/credulity/efficiency using session-level observations.

TABLE 3—AVERAGE BEHAVIORAL CHOICES AND CONSISTENT BELIEFS: ADDITIONAL TREATMENTS (LAST EIGHT SUPERGAMES)

Category	Measure	Partners	Chat	Transfer	Partners-R ₁	Partners-R ₂
Sender behavior	$\hat{\beta}(\text{Don't} \text{Bad})$	0.259 (0.031)	0.809 (0.029)	0.565 (0.039)	0.612 (0.034)	0.519 (0.036)
Consistent belief	$\hat{\mu}(\text{Good} \text{Invest})$	0.571 (0.010)	0.839 (0.021)	0.693 (0.020)	0.715 (0.018)	0.670 (0.017)
Receiver behavior	$\hat{\beta}(\text{Full} \text{Invest})$	0.251 (0.027)	0.830 (0.038)	0.466 (0.038)	0.601 (0.038)	0.536 (0.035)
	$\hat{\beta}(\text{None} \text{Don't})$	0.560 (0.065)	0.225 (0.035)	0.872 (0.028)	0.802 (0.036)	0.832 (0.035)
	with transfer:			0.596 (0.047)		
	$\hat{\beta}(\text{Partial} \text{Don't})$	0.388 (0.066)	0.729 (0.038)	0.124 (0.027)	0.179 (0.035)	0.151 (0.034)
	with transfer:			0.000 —		
Efficiency	$\hat{\text{Pr}}(\text{Full} \text{Good})$	0.284 (0.036)	0.861 (0.035)	0.521 (0.043)	0.621 (0.042)	0.586 (0.040)

Note: Standard errors (in parentheses) are derived from a bootstrap (of size 5,000) with supergame-level resampling, stratified by treatment.

the time in Partners and 86.1 percent in Chat. Qualitatively, the efficiency differences demonstrate that coordination failure is a first-order cause for the babbling play observed in Partners. While not perfectly coordinated on the efficient outcome, the large majority of rounds with preplay communication are efficient.

Beyond the efficiency increase, we also observe a change in the distributional outcomes in the bad state, further refining the conclusions. The modal response to a Don't Invest message in our Chat supergames is partial investment. The Don't Invest signal in the Chat treatment creates an empirically consistent belief on the bad state of 99 percent—so no investment is the clear, myopically rational response for receivers. However, Chat receivers almost double their selection of partial investment in response to this message, relative to Partners. This ends up being a direct effect of subjects explicitly coordinating on the payment of an information rent.

While the Chat treatment duplicates the supergame-level data on behavior collected in Partners, it additionally generates chat logs from subjects' conversations. An example of a chat interaction is the below exchange between a receiver subject (*R*) and a sender subject (*S*):³¹

R: I'll trust you until you lie and then it's [*Partial*] the whole way out.

S: Hey want to work together on this?

R: If you click [*Don't Invest*], I'll go [*Partial*] so we both get something

³¹The experiment had a left/middle/right (L/M/R) labeling for states, messages, and actions. We lightly edit the conversation with our economic labeling, indicating edits with square brackets. All other text/spelling mistakes are verbatim. In online Appendix F, we provide unedited chat logs from every single *Chat* interaction. Sequencing is by time stamp of hitting enter on the message.

S: I will tell you all the honest computer decisions if you never click [*None*]

S: instead when i mark [*Don't*] click [*Partial*]

S: deal?

R: no problem

R: deal

In the above conversation, both subjects near-simultaneously outline the information rent strategy, and this particular chat is from the very first supergame where they had the chat interface. Despite extensive realized failures to coordinate in the 12 prior supergames, both subjects recognize the information rent strategy as the way to coordinate.

While the above example chat was chosen for clarity, the discussion of the information rent strategy where partial investment is chosen after a bad state is revealed is entirely representative. To show this, we designed a coding protocol for chat logs with 23 questions (available in Table C2 of online Appendix C). Two research assistants (with no prior knowledge of our research hypotheses) independently coded each conversation according to the protocol. Examining the coded data and conservatively reporting averages only for data where the coders agree, we find that approximately three-quarters of chat exchanges have either the sender or receiver explicitly raise the information rent path of play. Explicit discussions of conditional punishments for deviations are rare, addressed in 10 percent of exchanges—as in the receiver's first statement in the example chat. When punishments are mentioned, they always refer to babbling reversion.³²

Examining the RA-coded data, in the first supergame where chat was made available, we find that 50 percent of the exchanges mention the information rent strategy. This figure grows to 86 percent by the final chat supergame.³³ While some subjects may learn and adopt the information rent strategy from a previous chat interaction, the fact that so many subjects discuss it at the first opportunity indicates that many are aware of the strategy but require communication to coordinate on it.³⁴

Using the coded chat data, our last check is that the efficiency gains in Chat are indeed driven by the information rent discussion.³⁵ Looking at the subset of

³²Supplementing the chat coding analysis, online Appendix E uses the Strategy Frequency Estimation Method (SFEM) of Dal Bó and Fréchette (2011) to provide estimates of the frequencies of use for a specified set of strategies. Approximately 90 percent of strategies in Chat are attributed to the information rent strategy with babbling reversion. While their dynamic play suggests babbling reversion *is* used to punish deviations, one reason subjects might avoid discussing the off-path punishments is as a rhetorical device to convince others, focusing instead on the benefits.

³³We focus here on supergames where our two coders agree, which is 22 of 24 for both the initial and last chat supergames.

³⁴There are no systematic differences in who initiates the strategy discussion: 47 percent of supergames mentioning the information rent strategy have the first mention by the sender, while 53 percent are by the receiver.

³⁵In the simplest information rent equilibrium, the payment of an information rent takes place during the same round in which the sender honestly reveals a bad state through a Partial response. But payment through the continuation is also possible when there is evidence of intertemporal transfers in other settings (Tavoni et al. 2011). Some subjects do discuss intertemporal transfers in their chats (for example, see the chats between Sender-86 and Receiver-140 in Session-14, Supergame-19 in online Appendix F). However, we observe little evidence in the data for actual intertemporal transfers. In the Chat treatment, there are 150 supergames where the relationship has thus far been both fully revealing and efficient until round t , where $t \in \{3, 4, 5\}$ and where in round $t + 1$ a Don't Invest message is received. Only 6 of these 150 supergames (4 percent) have the receiver select a higher rent payment through Full Investment, which would entail a transfer of \$3 to the sender.

partnerships that mention the information rent path in their chat, 91 percent have perfectly efficient play in the subsequent supergame. This figure contrasts with the 40 percent of partnerships that do not mention it.³⁶ As a summary of our Chat treatment, we have the following result.

FINDING 2: The evidence from observed behavior in Chat supergames and the chat logs is consistent with preplay communication successfully coordinating subjects on the information rent strategy, driving a substantial 57.7 percentage points increase in efficiency relative to Partners.

By subtracting the coordination problem from Partners, our Chat treatment indicates that the failure to coordinate on the mutually acceptable information rent strategy is driving the low efficiency outcomes. Absent the coordination problem, both senders and receivers show that they independently understand and jointly accede to the information rent outcome. Absent the tacit coordination problem over which dynamic strategy to choose, we find strong support support for the second alternative in Hypothesis 2.³⁷

B. Which Party Is Failing in Coordinating on the Information Rents Outcome?

From the Chat results, both senders and receivers are aware of the information rent strategy, but from Partners we know that subjects fail at tacitly coordinating on it. What then is the precise difficulty? One option is that senders find it difficult to convey their choice to reveal information to the receiver. For a sender honestly revealing the state, a bad-state realization in round one has some coordinative value, as the sender can use it to unambiguously separate from a dishonest sender who always says Invest. In a separate experimental treatment, we show that helping senders signal their honesty in the first round does not produce significant gains. As the result here is a null, we relegate the details and results from this treatment to an online Appendix for interested readers.³⁸

Another option for why tacit coordination fails is that receivers find it difficult to signal their intentions in paying an information rent. Through an additional treatment, we show that this does make a difference. The treatment modifies the repeated relationship Partners game by providing the receiver with an option to separate their action from the information rent payment. In this manipulation, which we call *Transfer*, the receiver makes two choices in response to messages. The first matches exactly the Partners' receiver actions: full, partial, or no investment. The second

³⁶For details and evidence for statistically significant differences driven by the information rent strategy, see Table C3 in online Appendix C, where we present a Tobit regression of the supergame efficiency on features of the preplay discussion.

³⁷While we view the Chat treatments as providing a coordination device, readers are also referred to the literature on the effect of promises (see Charness and Dufwenberg 2006 and citations thereof). The next treatment removes the ability to offer promises.

³⁸See online Appendix B. The treatment is identical to Partners except that senders decide on a message strategy each round, a message to send for each state without knowing the state realization. Receivers are shown the message for the realized state before they make their choice, but at the feedback stage, they also learn the complementary message.

involves a choice of a transfer amount to the sender: \$0 or \$1. Final payoffs are identical to the Partners treatment minus the transfer.³⁹

Because of the action manipulation, there are now two payoff-equivalent implementations of the simple information rent equilibrium. In both versions, the sender fully reveals the state, and the receiver selects *full investment* with no transfer when the good state is signaled. The difference is that when a bad state is signaled, the receiver can respond by combining their myopic best reply of no investment with a separate \$1 transfer (instead of the more ambiguous Partial choice). The bad-state payoffs from each implementation are therefore identical, with \$1 to the sender (\$0 from the action, +\$1 from the transfer) and \$2 for the receiver (\$3 from the action, -\$1 from the transfer). The fundamental difference in the Transfer modification to Partners is that receivers can separately convey their trust in the sender's message as well as their intention to pay a rent.⁴⁰

The third data column in Table 3 summarizes the aggregate behavior from the three Transfer sessions (46 subjects). Comparing the results to Partners, a first observation is that senders are more likely to honestly reveal. Truthful revelation in the bad state increases by 30 percentage points (pushing the empirically consistent belief moving above the critical threshold of $2/3$). The second observation is the increase in efficiency. The proportion of choices with full investment in the good state is 52 percent in our Transfer modification. This approximately doubles the efficiency from the Partners treatment, a significant increase in efficiency of almost 24 percentage points relative to Partners.

Consistent with the information rent strategy, there are very few transfers when the message is *Invest*. In the row below $\hat{\beta}(\text{Full}|\text{Invest})$, we provide the fraction of participants that chose full investment and a transfer, $\hat{\beta}(\text{Full}, \$1 \text{ transfer}|\text{Invest})$. Overall, just 5 percent of choices in response to the *Invest* message make the \$1 transfer alongside full investment, so few receivers are choosing to pay a rent when they believe the state is good. Transfers are also not made with partial investment choices (<0.01 after either message), nor alongside a no investment in response to *Invest*. The only situation where we find frequent transfer use is in response to a *Don't* message in combination with the myopically rational response of no investment.

Relative to Partners, the no-investment action in response to a signaled bad state is significantly more likely in Transfer (a 31 percentage point increase to 87 percent). However, the majority of these no-investment actions (68 percent) are coupled with a \$1 transfer to the sender that signaled *Don't Invest*. Pooling together the two alternative information rent responses (partial investment/no transfer and no investment/\$1 transfer), we find that 72 percent of receiver responses pay an information rent. This represents a 33 percentage point increase over *Partners* (significant at the 1 percent level), where modal behavior is now consistent with the

³⁹Formally, the action space for receivers is now $\hat{\mathcal{A}} = \{\text{Full}, \text{Partial}, \text{None}\} \times \{0, 1\}$, with a generic choice (a, x) being an action a and a transfer choice x where the round payoffs (in dollars) to the sender and receiver are $3u(a) + x$ and $3v(a, \theta) - x$, respectively. In addition, as the feedback at the end of the round also contains information on transfers, the game's history must also be adjusted.

⁴⁰While the addition of the transfer does expand the set of feasible payoffs for the repeated game (see Figure C1 in online Appendix C for a graphical depiction), the individually rational payoff sets are identical.

information rent equilibrium. The *Transfer* manipulation suggests that even a simple nudge toward the information rent strategy can create substantial efficiency gains.⁴¹ We now summarize the result as follows.

FINDING 3 (Explicit Rents): *A device that clearly signals the payment of an information rent separate from the action helps coordinate receivers on the information rent strategy. Efficiency relative to Partners increases by 23.7 percentage points.*

C. Robustness

The Chat treatment shows that subjects jointly agree on an information rent strategy when given an explicit coordination device. While not quite as powerful, the Transfer treatment shows that providing the ability to clearly convey an information rent through a separate explicit payment creates substantial efficiency gains over Partners. Taken together, these findings suggest that so long as the rent is unambiguously understood, we find strong support for the second part of Hypothesis 2. That is, where babbling punishments are used—and we find strong evidence for this when we examine the more detailed dynamic response—efficient outcomes require the payment of an information rent.

To stress test our findings in this final section, we introduce a final pair of treatments that modify the incentives within the Partners stage game so that (i) the efficient, fully extractive outcome (with no information rent) can be supported with babbling reversion and (ii) the babbling partial action does not implement a transfer to the sender relative to no investment. This second feature means that we remove any ambiguity in a receiver's choice of partial investment. The hypothesis here is that if we remove the need (and ability) to pay an information rent, we will also remove the main coordination hurdle.⁴² Relative to Partners, we will show that both modifications produce comparable efficiency gains to those found in the Transfer treatment.

We designed two treatments that change the sender's incentives in the Partners game using the payoffs Robustness₁ and Robustness₂ provided in Table 1, panel A.⁴³ The first modification, *Partners-R*₁, increases the sender's payoff from no investment from \$0 to \$1. The second, *Partners-R*₂, reduces the sender payoff when the receiver selects partial investment from \$1 to \$0. In both treatments, there is the same babbling stage game Nash equilibrium where the receiver selects partial investment. Both modifications allow for the Full Extraction outcome as an equilibrium of the repeated game, supported with babbling reversion.⁴⁴ In *Partners-R*₁ the punishment payoffs are the same as in Partners, but the on-path payment from full

⁴¹ Our focus in the body of the paper is on simple aggregates, as these best illustrate the core finding. A more detailed analysis of the supergame behavior indicates that the information rent is supported by a history-dependent babbling trigger (see online Appendix E).

⁴² These treatments are useful as robustness tests, but they capture an extreme in which babbling supports Full Extraction. As indicated in the theory section, the more general situation is one in which information rents can be supported with reversion to babbling, which is captured by our Partners treatment.

⁴³ Figure C2 in online Appendix C provides a characterization of the individually rational payoff pairs in the two treatments.

⁴⁴ An effect of both changes is that the static Nash (babbling) and the worst-case individually rational (minmax) payoffs now coincide.

extraction leaves the sender with \$1 instead of \$0 when the receiver's choice is no investment. For Partners-R₂ the babbling punishment becomes harsher relative to Partners, leaving the sender with a \$0 payoff in comparison to the \$1 they received in the original version. Overall, both payoff changes have the effect of making Full Extraction the unique efficient stage game outcome supportable in an equilibrium of the repeated game.⁴⁵ While the Partial action in the Partners treatment can be used to implement efficiency, in both Partners-R manipulations, it clearly indicates the implementation of a nonefficient outcome.

For each Partners-R treatment, we conducted three sessions with a total of 46 subjects, with results provided in the last two columns of Table 3. Both modifications show significant increases in revealing, efficient, and fully extractive play. Focusing on our efficiency measure given by $\Pr(a = \text{Full} | \theta = \text{Good})$, both Partner-R treatments more than double the Partners efficiency level when we look at aggregate behavior.⁴⁶

FINDING 4 (Robustness): *Shifting the Partners game payoffs to remove the requirement for an information rent to support revelation (via babbling reversion), behavior in both Partners-R treatments increases the coordination on efficient, revealing play by approximately 30 percentage points relative to Partners.*

V. Conclusion

Our paper experimentally examines a repeated cheap talk setting, a setting where independent interpretation of two separate literatures on behavior in one-shot cheap talk and repeated prisoner's dilemma (RPD) games would predict substantial improvements over the Bayesian Nash stage game prediction. However, our main results instead show a robust coordination on inefficient babbling outcomes, with only economically small increases from repeated play. Through a number of small experimental modifications, we diagnose the cause: coordination failure over how efficiency gains are distributed. Where the standard RPD game focuses the coordination problem solely on the efficiency margin, a feature of repeated cheap talk—and a generic feature in many repeated games—is a requirement to additionally resolve distributional issues. If only simple dynamic punishments can be coordinated on, then efficient play in standard repeated cheap talk will require the payment of an information rent to the sender. The main contribution of our paper is to show that tacit coordination on the payment of this rent does not arise without some help.

While a strong coordination device (free-form preplay communication) essentially solves the problem, we show that even a weak coordination device focused on making the payment of an information rent less ambiguous can produce substantial

⁴⁵ All efficient outcomes require the receiver to choose (or mix) between Full and None when the state is Bad (Partial is Pareto dominated in both states in our robustness games). However, at $\delta = 3/4$, Full cannot be chosen with positive probability at a revealed bad state in any stationary equilibrium of the game. Without the possibility of paying a rent intertemporally, full extraction is the unique efficient stage game outcome supportable as a PBE.

⁴⁶ The Partners-R₁ treatment (where we lower senders' payoffs relative to Partners) performs significantly better than the Partners-R₂ in round one cooperation. However, looking across all supergame rounds, receiver behavior and final efficiency are not significantly different from one another, even at the 20 percent level. Senders' increased revelation in Partners-R₁ is significantly greater, at the 10 percent level.

gains. Enriching the action space to allow the receivers to make a separate side payment to the sender, we generate approximately half of the efficiency gains from preplay communication. Where tacit coordination on efficient dynamic strategies fails in the original environment, even a relatively weak device to make the information rent payment explicit can substantially alter final outcomes.

Finally, in two robustness treatments, we hold constant the original repeated relationship setting but modify the sender's payoffs to remove the requirement for an information rent payment to support efficient play (via babbling reversion). In each modification, we find efficient, fully extractive play as the modal outcome.

While our results point toward the importance of coordination problems in repeated settings (and some potential remedies), there are many questions left for future research to clarify. One question that jumps out is the extent to which the availability of punishments play an empirical role in equilibrium selection. Across all of our experiments, observed behaviors point toward the idea that rather than the individually rational payoff set, it is the set of payoffs supported by Nash reversion that is important for prediction. Though outside the scope of this paper, the next step is to understand the limits of this finding and how it generalizes across contexts, as it could be extremely useful in reducing the indeterminacy caused by more permissive folk theorems using minmax punishments.

REFERENCES

- Abeler, Johannes, Daniele Nosenzo, and Collin Raymond.** 2016. "Preferences for Truth-Telling." *Econometrica* 87 (4): 1115–53.
- Abreu, Dilip, and Yuliy Sannikov.** 2014. "An Algorithm for Two-Player Repeated Games with Perfect Monitoring." *Theoretical Economics* 9: 313–38.
- Arechar, Antonio A., Anna Dreber, Drew Fudenberg, and David G. Rand.** 2017. "I'm Just a Soul Whose Intentions Are Good": The Role of Communication in Noisy Repeated Games." *Games and Economic Behavior* 104: 726–43.
- Baron, David P., and David Besanko.** 1984. "Regulation and Information in a Continuing Relationship." *Information Economics and Policy* 1 (3): 267–302.
- Blume, Andreas, Douglas V. DeJong, Yong-Gwan Kim, and Geoffrey B. Sprinkle.** 1998. "Experimental Evidence on the Evolution of Meaning of Messages in Sender-Receiver Games." *American Economic Review* 88 (5): 1323–40.
- Brown, Martin, Armin Falk, and Ernst Fehr.** 2004. "Relational Contracts and the Nature of Market Interactions." *Econometrica* 72 (3): 747–80.
- Cai, Hongbin, and Joseph Tao-Yi Wang.** 2006. "Overcommunication in Strategic Information Transmission Games." *Games and Economic Behavior* 56 (1): 7–36.
- Camera, Gabriele, and Marco Casari.** 2009. "Cooperation among Strangers under the Shadow of the Future." *American Economic Review* 99 (3): 979–1005.
- Charness, Gary, and Martin Dufwenberg.** 2006. "Promises and Partnership." *Econometrica* 74 (6): 1579–1601.
- Cooper, David J., and John H. Kagel.** 2005. "Are Two Heads Better than One? Team versus Individual Play in Signaling Games." *American Economic Review* 95 (3): 477–509.
- Crawford, Vincent P., and Joel Sobel.** 1982. "Strategic Information Transmission." *Econometrica* 50 (6): 1431–51.
- Dal Bó, Pedro, and Guillaume R. Fréchet.** 2011. "The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence." *American Economic Review* 101 (1): 411–29.
- Dal Bó, Pedro, and Guillaume R. Fréchet.** 2018. "On the Determinants of Cooperation in Infinitely Repeated Games: A Survey." *Journal of Economic Literature* 56 (1): 60–114.
- Dickhaut, John W., Kevin A. McCabe, and Arijit Mukherji.** 1995. "An Experimental Study of Strategic Information Transmission." *Economic Theory* 6 (3): 389–403.

- Dreber, Anna, David G. Rand, Drew Fudenberg, and Martin A. Nowak.** 2008. "Winners Don't Punish." *Nature* 452: 348–51.
- Duffy, John, and Jack Ochs.** 2009. "Cooperative Behavior and the Frequency of Social Interaction." *Games and Economic Behavior* 66 (2): 785–812.
- Engle-Warnick, Jim, and Robert L. Slonim.** 2004. "The Evolution of Strategies in a Repeated Trust Game." *Journal of Economic Behavior and Organization* 55 (4): 553–73.
- Engle-Warnick, Jim, and Robert L. Slonim.** 2006. "Inferring Repeated-Game Strategies from Actions: Evidence from Trust Game Experiments." *Economic Theory* 28 (3): 603–32.
- Ettinger, David, and Philippe Jehiel.** 2016. "An Experiment on Deception, Reputation and Trust." <https://philippe-jehiel.enpc.fr/wp-content/uploads/sites/2/2019/01/expedeception.pdf>.
- Fischbacher, Urs.** 2007. "z-Tree: Zurich Toolbox for Ready-Made Economic Experiments." *Experimental Economics* 10 (2): 171–78.
- Fischbacher, Urs, and Franziska Föllmi-Heusi.** 2013. "Lies in Disguise—An Experimental Study on Cheating." *Journal of the European Economic Association* 11 (3): 525–47.
- Fréchette, Guillaume R., and Emanuel Vespa.** 2017. "The Determinants of Voting in Multilateral Bargaining Games." *Journal of the Economic Science Association* 3 (1): 26–43.
- Fudenberg, Drew, and Eric Maskin.** 1986. "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information." *Econometrica* 54 (3): 533–54.
- Fudenberg, Drew, David G. Rand, and Anna Dreber.** 2012. "Slow to Anger and Fast to Forgive: Cooperation in an Uncertain World." *American Economic Review* 102 (2): 720–49.
- Gneezy, Uri.** 2005. "Deception: The Role of Consequences." *American Economic Review* 95 (1): 384–94.
- Goeree, Jacob K., and Charles A. Holt.** 2000. "Asymmetric Inequality Aversion and Noisy Behavior in Alternating-Offer Bargaining Games." *European Economic Review* 44 (4–6): 1079–89.
- Golosov, Mikhail, Vasiliki Skreta, Aleh Tsyvinski, and Andrea Wilson.** 2014. "Dynamic Strategic Information Transmission." *Journal of Economic Theory* 151: 304–41.
- Güth, Werner, and Martin G. Kocher.** 2014. "More than Thirty Years of Ultimatum Bargaining Experiments: Motives, Variations, and a Survey of the Recent Literature." *Journal of Economic Behavior and Organization* 108: 396–409.
- Hurkens, Sjaak, and Navin Kartik.** 2009. "Would I Lie to You? On Social Preferences and Lying Aversion." *Experimental Economics* 12 (2): 180–92.
- Kagel, John H., and Peter McGee.** 2016. "Team versus Individual Play in Finitely Repeated Prisoner Dilemma Games." *American Economic Journal: Microeconomics* 8 (2): 253–76.
- Kartal, Melis, Wieland Müller, and James Tremewan.** 2017. "Building Trust: The Costs and Benefits of Gradualism." <https://pdfs.semanticscholar.org/5f24/6a1f181b7b9248dbb2e1e711208558785253.pdf>.
- Katzwer, Richard.** 2013. "rgsolve for Java." Unpublished.
- Laffont, Jean-Jacques, and Jean Tirole.** 1988. "The Dynamics of Incentive Contracts." *Econometrica* 56 (5): 1153–75.
- Lai, Ernest K., Wooyoung Lim, and Joseph Tao-yi Wang.** 2015. "An Experimental Analysis of Multi-dimensional Cheap Talk." *Games and Economic Behavior* 91: 114–44.
- Le Quement, Mark T., and Amrith Patel.** 2018. "Communication as Gift-Exchange." University of East Anglia Working Paper 2018-06.
- Mathevet, Laurent, David Pearce, and Ennio Stacchetti.** 2019. "Reputation and Information Design." <http://www.laurentmathevet.com/wp-content/uploads/2019/02/MPS.pdf>.
- Renault, Jérôme, Eilon Solan, and Nicolas Vieille.** 2013. "Dynamic Sender-Receiver Games." *Journal of Economic Theory* 148 (2): 502–34.
- Roth, Alvin E.** 1995. "Chapter 4—Bargaining Experiments." In *Handbook of Experimental Economics*, edited by John H. Kagel and Alvin E. Roth, 253–348. Princeton, NJ: Princeton University Press.
- Sobel, Joel.** 1985. "A Theory of Credibility." *Review of Economic Studies* 52 (4): 557–73.
- Sutter, Matthias.** 2009. "Deception through Telling the Truth?! Experimental Evidence from Individuals and Teams." *Economic Journal* 119 (534): 47–60.
- Tavoni, Alessandro, Astrid Dannenberg, Giorgos Kallis, and Andreas Löschel.** 2011. "Inequality, Communication, and the Avoidance of Disastrous Climate Change in a Public Goods Game." *Proceedings of the National Academy of Sciences* 108 (29): 11825–29.
- Vespa, Emanuel, and Alistair J. Wilson.** 2016. "Communication with Multiple Senders: An Experiment." *Quantitative Economics* 7 (1): 1–36.
- Vespa, Emanuel, and Alistair J. Wilson.** 2017. "Experimenting with Equilibrium Selection in Dynamic Games." http://www.pitt.edu/~alistair/papers/MPE_June2017.pdf.

- Vespa, Emanuel, and Alistair J. Wilson.** 2019. "Experimenting with the Transition Rule in Dynamic Games." *Quantitative Economics* 10 (4): 1825–49.
- Wang, Joseph Tao-yi, Michael Spezio, and Colin F. Camerer.** 2010. "Pinocchio's Pupil: Using Eyetracking and Pupil Dilation to Understand Truth Telling and Deception in Sender-Receiver Games." *American Economic Review* 100 (3): 984–1007.
- Wilson, Alistair J., and Emanuel Vespa.** 2020. "Replication Data for: Information Transmission under the Shadow of the Future: An Experiment." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. <https://doi.org/10.3886/E115922V1>.
- Wilson, Alistair J., and Hong Wu.** 2017. "At-Will Relationships: How an Option to Walk Away Affects Cooperation and Efficiency." *Games and Economic Behavior* 102: 487–507.