INFORMATION TRANSMISSION UNDER THE SHADOW OF THE FUTURE: AN EXPERIMENT

ALISTAIR J. WILSON AND EMANUEL VESPA

ABSTRACT. We experimentally examine how information transmission functions in an ongoing relationship. Where the one-shot literature documents substantial over-communication and preferences for honesty, in our repeated setting the outcomes 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 the inefficient outcomes we find are driven by a coordination failure. More specifically, we show that in order to sustain honest revelation, the focal dynamic punishments in this environment require a coordination on the payment of "information rents."

"'tis like the breath of an unfee'd lawyer. You gave me nothing for't. Can you make no use of nothing[?]"

– *King Lear*, (1.4.642–44, 1.4.660–62)

Our paper examines the intersection of two key areas of current economic research: the degree to which information is shared by interested parties and how repeated interaction in long-run relationships can open up the possibility for more-efficient outcomes. Our main question focuses on whether repeated interaction can produce efficiency-improvements to the transmission of information. Our findings for repeated interaction indicate a qualified positive, where the main motive behind greater information transmission when we observe it are strategic: information is shared only when there is compensation. However, repeated interaction on its own is not sufficient to generate an efficiency gain; overcoming coordination hurdles on how to share efficiency gains are critical to successful outcomes.

From a theoretical perspective, repeated interactions clearly offer the potential for gains with information transmission. While full revelation is not an equilibrium in one-shot settings such as

Date: November, 2016.
We would like to thank: Ted Bergstrom, Andreas Blume, Gabriele Camera, Katie and Lucas Coffman, John Duffy, Rod Garratt, PJ Healy, Navin Kartik, Andrew Kloosterman, John Kagel, Tom Palfrey, Joel Sobel, Lise Vesterlund, Sevgi Yuksel, as well as seminar audiences at Ohio State, Penn State, Arizona, Pittsburgh, UC Berkeley, UC Irvine (SWEBE), UC Riverside, UC Santa Cruz, UC San Diego, UC Santa Barbara, UCL, Virginia, SITE and the international meeting of the ESA in Jerusalem. 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; this order conveys no information on contribution which was entirely equal. Wilson: alistair@pitt.edu; Vespa: vespa@ucsb.edu.

Crawford and Sobel (1982), it can be supported if parties interact repeatedly. Moreover, two current literatures taken independently provide an optimistic prior. First, the experimental literature on one-shot information-transmission games documents over-transmission relative to the theoretical prediction (see Abeler et al., 2016, for a meta-study). Second, a literature examining repeated interaction in prisoners' dilemmas (see Dal Bó and Fréchette, 2016, for a meta-study) shows that repeated interactions can improve efficiency through the use of dynamic strategies. However, caution may be warranted when extending these two results to repeated cheap talk. With respect to the first finding, informed parties may be discouraged from honest revelation if confronted with a lack of reciprocity from decision makers who use the provided information solely to their own advantage. This filters into a problem for the second finding: the coordination problem in this environment is *not* a trivial extension of that in a repeated prisoners' dilemma. In a repeated cheap talk environment many efficient equilibria can be supported depending on how the informed and uninformed participants divide the efficiency gains from information transmission. A subtle trade-off emerges for coordination: the closer the uninformed party's payoff is to their first best, the more-severe the required off-path punishments to discipline informed parties that lie.

In our experiments we use a simple sender-receiver environment with uncertainty over a payoff-relevant state. In our 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 (DM, the receiver, he), who chooses a level of investment. We restrict the sender to choose 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. But, in addition to full and no investment, the DM can also select an intermediate *partial* investment level, which is the best response to his initial prior on the opportunity.

The incentives create a clear strategic tension. While there is mutual benefit from full investment for a good opportunity, the preferences of the two parties are completely misaligned for a bad one. In one-shot interactions the wedge between honest revelation and self-interest for low-quality investments leads to substantial inefficiencies in equilibrium. For our particular setting the unique equilibrium prediction is that the DM chooses his uninformed best response (here partial investment) at all quality levels. This uninformative "babbling" outcome is ex post inefficient, as full investment for good opportunities makes both expert and DM better off.

In expectation this environment is similar to a prisoner's dilemma (PD) game: a Pareto efficient outcome that cannot be achieved in one-shot interactions as each party benefits from a deviation. (The sender dishonestly claiming a good state, the DM reducing their investment for a revealed bad state.) Given the parallels to a prisoner's dilemma, our main paper asks an obvious question:

to what extent can repetition and feedback on past behavior increase efficiency? Theoretically, so long as each party values the future enough, more-efficient equilibria can be constructed if past play is observed. Much like an indefinitely repeated PD, history-dependent strategies can implement efficient outcomes (and many others) in equilibrium. However, unlike the PD, a novel coordination tradeoff emerges between: (i) harshness of the punishments; and (ii) how the surplus generated from revelation is distributed.

From the point of view of the DM, the first-best outcome is full-revelation of the investment quality by the expert, and maximizing his own payoff. The path of play for full extraction is cognitively simple for the DM, choosing his most-preferred outcome given the revealed quality: full investment for good opportunities, no investment for bad. Coordination on this path is also cognitively simple for the expert (honestly reveal). However, as none of the gains are shared, the net incentive to truthfully reveal bad opportunities must be driven by the threat of punishment. While possible in equilibrium, if DMs hope to support full extraction, they must use complex punishments that are (at least temporarily) painful to both parties. Moreover, such punishments must eventually recoordinate on efficient play.

In contrast to full extraction, simpler punishments can be used to support information revelation if the DM shares the efficiency gains from revelation with the expert—what we will refer to as an "information rent." In particular, given a large-enough information rent, an intuitively appealing dynamic punishment can be used to support efficient play in equilibrium: reversion to the uninformative babbling outcome. In the theoretical section of the paper, we show that the choice between coordinating on harsher punishments than the static Nash outcome, or coordination on the payment of information rents, is not unique to our experimental setting. Instead, under typical preference assumptions in these games, information rents are a necessary component of efficient play supported by Nash-triggers, even for future discount rates arbitrarily close to one.

While the literature on repeated play has found evidence for coordination on more-lenient punishments than Nash reversion (for instance, see Fudenberg et al., 2010), there is little evidence for coordination on harsher punishments. This behavioral observation motivates a refinement for folk theorems: a proscription on punishments harsher than the worst stage-game Nash equilibrium. In our sender-receiver environment, this simple refinement has two implications: First, if punishments can be no harsher than Nash, whenever efficient play does emerge it must necessarily involve the payment of an information rent. Second, as a parallel to the first implication, solving the coordination problem to achieve efficient outcomes is much tougher than standard repeated PD games, as it requires tacit agreement both on revelation and division.

Our paper's experimental design first focuses on measuring the effects of repeated play in our cheap-talk environment. A *Partners* treatment, where matchings are fixed within all periods of a supergame, is compared to a *Strangers* control in which participants are matched with a new subject every period. If repeated interaction selects equilibria with greater information transfer, we should observe large efficiency differences between these two treatments. Instead, we find that the large majority of subjects in both treatments coordinate on the uninformative babbling outcome. While we do document differences in the expected direction—more information is transmitted in our *Partners* treatment—the effects are economically small. While this result is surprising when set against the one-shot cheap-talk and infinite-horizon PD literatures, it is less so when viewed through the lens of the outlined coordination problems.

To investigate whether the babbling outcomes in our repeated environment are driven by an inability to coordinate on efficient play, we conduct an additional treatment which alleviates this issue. Though identical to our *Partners* treatment for the initial supergames, in the second part of our *Chat* treatments we let participants freely communicate before the supergame begins. Behavior is not distinguishable from *Partners* before the coordination device is provided, but large differences emerge from the very first supergame where they have access to it. While the comparable late-stage *Partners*' supergames are 24 percent efficient, efficiency is 86 percent in the *Chat* treatment. Moreover, a detailed examination of the chat exchanges clearly indicates that the large-majority of subjects are exactly coordinated on an efficient information-rents outcome.[1] That this outcome emerges at the very first opportunity to explicitly coordinate points to subjects' clear understandings of how to (jointly) solve their coordination problem.

The *Chat* treatment results indicate that where the coordination problem is relaxed, subjects pick out paths of play where information rents are provided to the expert. Our last two treatments investigate if it is possible to achieve efficient outcomes without pre-play communication by relaxing the coordination problem in alternative ways. Coordination in our setting can be thought of as having two components. The first is on whether information is revealed or not, which determines the coordination on efficient or inefficient outcomes. The second component only matters when information is fully revealed, as subjects need to solve a distributional issues and coordinate on one of the many possible efficient equilibria. Our first manipulation provides a device to reduce the difficulties coordinating on fully revealing outcomes. Here our results show no effect relative to *Partners*, with behavior again best characterized as uninformative. Our second manipulation instead provides a coordination device that helps address the distributional concerns key to the information rent. In this second manipulation we observe a significant increase in efficiency relative to *Partners*, attaining approximately 60 percent of the efficiency gains in the *Chat* treatment. Solving the coordination problem over the distribution given truthful revelation is critical for successful outcomes in this environment.

The evidence from our experimental treatments points to the importance of information rents, an idea that goes back to an older literature on regulating a monopolist's price (see Baron and Besanko,

[1]In contrast, full extraction by the DM, which requires harsh supporting punishments is hardly ever coordinated upon.

1984; Laffont and Tirole, 1988). To retain incentives for honest revelation on profitability, that literature shows that a monopolist must be given some positive profits along the path, as otherwise they have no incentive to provide information against their own interests. While our infinite-horizon setting admits the equilibrium possibility for full revelation without a rent (through harsh punishments) the data indicates that such outcomes are not selected. Once a relationship is deceptive, babbling is the absorbing outcome across our treatments. Where dynamic strategies of the game are restricted to such punishments, information rents become necessary for efficient play.

Our paper's conclusions mirror the quote that we open the paper with, honest advice must be paid for. But the context of that quote is important here too, and speaks to the difficulty of coordination. When we are being advised, we can focus too myopically on the specific advice given, and forget that our subsequent choices have implications for those providing expertise. Where we fail to reward those providing honest information today, one outcome is to degrade the extent to which we get honest advice tomorrow.

We next present a brief review of the literature. In Section 2 we outline our main treatment and control, along with the theoretical predictions. Section 3 contains the results from these two treatments. In Section 4 we introduce our *Chat* treatment and the results. Section 5 follows with a discussion of two weaker coordination-device treatments. Finally, in Section 6 we discuss our results and conclude.

1.1. **Literature.** Per the introductory paragraph, the paper is most clearly related to two literatures: information transmission and repeated games. For the first, the theory 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 finding here is that full revelation is ruled out (for the one-shot game) with at best partial revelation in equilibrium. The model has been examined in experimental studies, where the main results are that more information is exchanged than predicted by the theory.[2]

Outside of the Crawford and Sobel framework, Gneezy (2005) compares the willingness of an informed sender to signal the payoff-maximizing action to a receiver.[3] Despite large monetary losses from honest revelation, the paper documents substantial information transfer. Comparing the sender behavior to decisions in a matched dictator game, the paper concludes there is a preference for honesty.[4] Further evidence for lying-aversion is given in Fischbacher and Föllmi-Heusi (2013), introducing a die-rolling task where the subjects are paid based on their own private roll of a

[2]See Dickhaut et al. (1995); Cai and Wang (2006); Wang et al. (2010). Also see Blume et al. (1998) for an examinations of the evolution of message meaning.

[3]Receivers here are ignorant of the game form, only knowing that they have two available actions.

[4]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.

die (see Abeler et al., 2016, for a meta-study of such experiments). While the die outcomes are private, lying can be detected in aggregate by comparing the distributions of reports to the known generating process. Such deception environments are most comparable to the bad state in our game—with a direct tension between truthful reporting and self-interest. However, the predicted effects of lying aversion are toward increased efficiency, in opposition to our results.

Our paper is also related to the literature on conditional cooperation in repeated games. Much of the theoretical literature here focuses on documenting the generality and breadth of the indeterminacy of prediction via folk theorems.[5] While recent theoretical work examines dynamic sender-receiver environments with persistent state variables (see Renault et al., 2013; Golosov et al., 2014), our game has a much simpler separable structure.[6] With perfect-monitoring, standard repeated-game constructions can be used to show the existence of more-efficient outcomes. Though we do not make a technical contribution to this literature, we do outline a practically motivated refinement for these games that reduces the theoretical indeterminacy: punishments should be simple.

A robust experimental literature on behavior in repeated PD-games has examined predictable features of selection. Surveying the literature and conducting a meta-analysis Dal Bó and Fréchette (2016) show that cooperation being supportable as an equilibrium outcome is necessary.[7] But more than this, they show that the difficulty of the coordination problem (here measuring the efficiency gains available from cooperation measured relative to the temptations and risks from a deviations) are predictive of selection. Further, the evidence from the meta-study suggests that selected punishments are frequently less harsh than grim-triggers, with Tit-for-Tat a frequently detected strategy. In a similar vein, punishment phases in imperfect public monitoring games are typically forgiving (see Fudenberg et al., 2010) and so less harsh than the static Nash outcome. In dynamic game settings with state-variables Vespa and Wilson (2015) indicates that more-cooperative outcomes than the history-independent Markov perfect equilibria (MPE) can be sustained, but that the most-frequent punishments are reversions to the MPE. While some papers have examined the selection of more costly punishments than Nash reversion, none find that this selection is useful in supporting higher payoffs.[8] 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.[9]

[5]See Fudenberg and Maskin (1986); Fudenberg et al. (1994); Dutta (1995)

[6]For an experiment here see Ettinger and Jehiel (2015), which examines a repeated sender-receiver game, where the sender has a persistent type, mirroring the Sobel (1985) model of credibility.

[7]For the effect of pre-play communication in repeated PD-games see Arechar et al. (2016) and references therein. Also see the literature on team communication, in particular Cooper and Kagel (2005) and Kagel and McGee (2016).

[8]Dreber et al. (2008) find that where costly punishments are used they are unsuccessful at supporting better outcomes; similarly, Wilson and Wu (2016) find subjects employing a dominated strategy (termination) as a costly punishment, but without an observed increase in cooperation.

[9]Outside of the infinite-horizon setting, Brown et al. (2004) show that gains over the inefficient equilibrium prediction need to be shared for efficient outcomes to emerge in their relational contracting setting. In an infinite-horizon setting, these results can be rationalized with simple punishments based on termination and rematching. If gains are not shared

2. DESIGN: THEORY AND HYPOTHESES

In this section we present the basic environment that we study in the laboratory, a streamlined version of a repeated cheap-talk game. One alternative would be to directly use a Crawford and Sobel (1982) stage-game. Instead, we simplify the stage-game for two reasons. First, with a simpler stage game that retains the main tensions we reduce the complexity naturally added by repetition. Second, the analysis and interpretation of the data will be more direct without a multiplicity of stage-game Nash equilibria.

The stage game of the infinite-horizon sender-receiver environment that we study is depicted in Figure 1. The game unfolds as follows: (i) Nature chooses a state of the world $\theta \in \{G(\text{ood}), B(\text{ad})\}$, with equal probability; (ii) the *Sender* player perfectly observes the state θ and chooses a message $m \in \{\text{Invest}, \text{Don't Invest}\}$; (iii) the *Receiver* player 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 period (θ, m, a) becomes common knowledge.[10]

Payoffs for the game are as follows: The sender's payoff is given by

$$u(a) = \begin{cases} 1 & \text{if } a = \text{Full}, \\ u_0 & \text{if } a = \text{Partial}, \\ 0 & \text{if } a = \text{None}, \end{cases}$$

while the receiver's payoff is

$$v(a, \theta) = \begin{cases} 1 & \text{if } (a, \theta) \in \{\text{(Full, Good),(None,Bad)}\}, \\ v_0 & \text{if } a = \text{Partial}, \\ 0 & \text{otherwise.} \end{cases}$$

The parameters $u_0 \in (0, 1)$ and $v_0 > \frac{1}{2}$ represent how the two participants rank partial investment.

Conditional on the *Good* state, the two players share common interests: the investment actions are Pareto ranked (from best to worst) by *Full* to *Partial* to *None*. However, conditional on the *Bad* state the Sender and Receiver have diametrically opposed preferences over the selected actions. In particular, when θ =*Bad*, the *Full* investment choice distributes the entire amount to the sender, while *None* distributes everything to the receiver. Finally, the preference assumption $v_0 > 1/2$ for

in one relationship, but are in the larger population, the simple punishment is to terminate the relationship and rematch, with shared gains a necessary implication.

[10]In 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 E*.

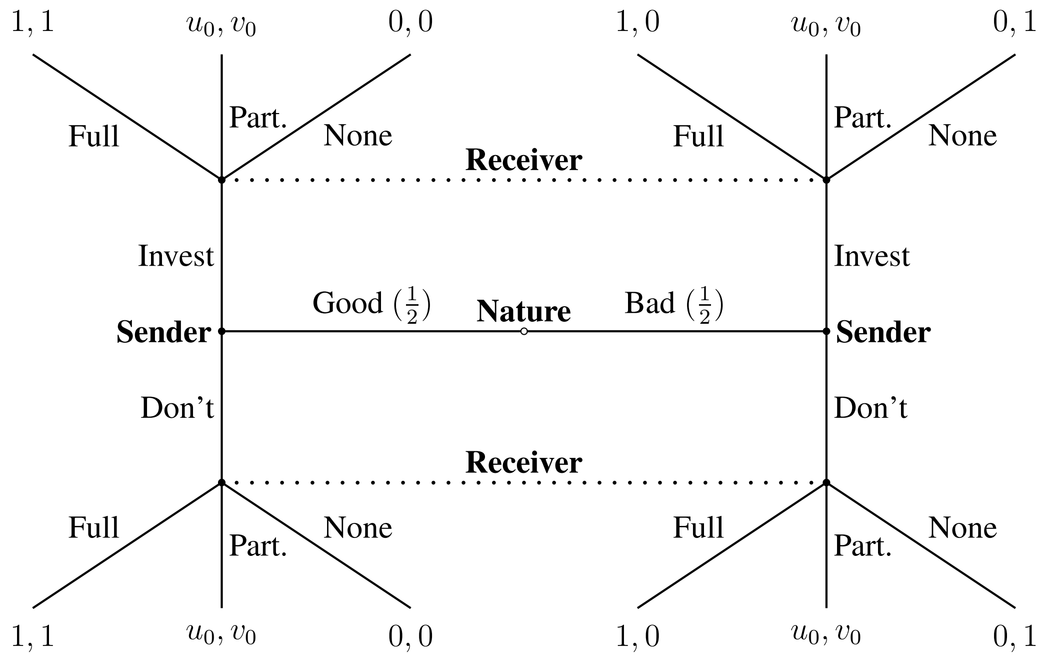

FIGURE 1. Stage Game

the receiver is chosen so that ex ante the *Partial* investment action is preferred to either full or no investment.[11]

The informed sender's chosen signal m affects neither players' payoff directly, and so this is a cheap-talk environment. Moreover the sender's payoff here is state independent, affected by the chosen action only. In terms of theoretical predictions for the stage game, this means that all perfect Bayesian equilibrium (PBE) involve the receiver choosing $a = \textit{Partial}$ with certainty.[12]

2.1. **Experimental Design.** Our experimental design will examine a repeated (infinite-horizon) version of the stage game given in Figure 1, which we will refer to as a supergame. We choose $u_0 = 1/3$ and $v_0 = 2/3$ so that the payoffs in the bad state are constant sum for all actions, and in that way it is a purely distributional state.

[11]This is under the assumption of risk-neutrality over the payoffs. In order to move away from the babbling prediction the receiver must be risk-loving to a very strong degree (a CRRA coefficient -0.71 or below). For risk-neutral preferences a posterior belief on either state above v_0 is required for partial investment not to be the best response, with this requirement increasing in their risk aversion (effectively pushing up the value of v_0).

[12]The 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. For a stage game with only two choices: (Full and None) the receiver can implement babbling by ignoring the message and selecting Full, None, or a mixture. In all cases the expected payoff is $1/2$.

After each completed period in the supergame we induce a $\delta = 3/4$ probability that the supergame continues into another period, and a $(1-\delta)$ probability that the supergame stops. From the point of view of a risk-neutral subject in period one, the expected payoffs for the sequence of possible states/messages/actions for the sender and receiver, respectively, is proportional to:

$$U\left(\{m^t(\theta), a^t(m)\}_{t=1}^{\infty}\right) = \quad (1-\delta)\sum_{t=1}^{\infty} \delta^{t-1}\mathbb{E}_{\theta} u\left(a^t\left(m^t\left(\theta\right)\right)\right),$$

$$V\left(\{m^t(\theta), a^t(m)\}_{t=1}^{\infty}\right) = \quad (1-\delta)\sum_{t=1}^{\infty} \delta^{t-1}\mathbb{E}_{\theta} v\left(a^t\left(m^t\left(\theta\right)\right), \theta^t\right).$$

We will initially examine two treatments: In the first, which we call our ***Partners*** treatment, the sender-receiver pair is fixed across the supergame. A sender participant i, and a receiver participant j are therefore paired together in all supergame periods. A sender in period $T > 1$ observes the history $h_i^T = \left\{\theta^T, \{\theta^t, m_i^t, a_j^t\}_{t=1}^{T-1}\right\}$ when it is their turn to move, while the receiver observes $h_j^T = \left\{m_i^T, \{\theta^t, m^t, a^t\}_{t=1}^{T-1}\right\}$.

In our second treatment, which we call ***Strangers***, the sender-receiver pairings in each period are random. A particular sender i is anonymously matched with a random sequence of receivers, $J(t)$. Similarly each receiver j is matched to a random sequence of senders, $I(t)$. The sender's observed history is therefore $h_i^T = \left\{\theta^T, \left\{\theta^t, m_i^t, a_{J(t)}^t\right\}_{t=1}^{T-1}\right\}$ where they have no knowledge of the particular receivers they have interacted with. Similarly, the receiver's observable history in each period is $h_j^T = \left\{m_{J(T)}^T, \left\{\theta^t, m_{J(t)}^t, a_i^t\right\}_{t=1}^{T-1}\right\}$, where the receiver again does not know which senders they have previously interacted with.[13]

2.2. **Strangers Game.** Given the random, anonymous matching, we treat each supergame in our *Strangers* treatment as a sequence of one-shot games. Though the instructions, supergame lengths, and potential payoffs are the same as our *Partners* treatment, subjects are unable to tie the previous history to the presently matched participant. The shadow of the future can not work if future behavior is independent of history. In principle, were δ large enough, more-efficient equilibria *are* possible in the *Strangers* treatment, through "contagion" constructions. We will not focus on such constructions for two reasons: First, experimental examination of contagion equilibria show that they lack predictive power when they do exist.[14] Second, given our choice of $\delta = 3/4$ and the particular random-matching protocol we use, such outcomes are not equilibria of the induced session-level supergame.

[13]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 supergame periods: so $\Pr\{I(t) = I(t+1)\}$=$\Pr\{J(t) = J(t+1)\} = 0$ for all t.

[14]Duffy and Ochs (2009) finds no support for contagion equilibrium predictions. Camera and Casari (2009) show 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 is 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.

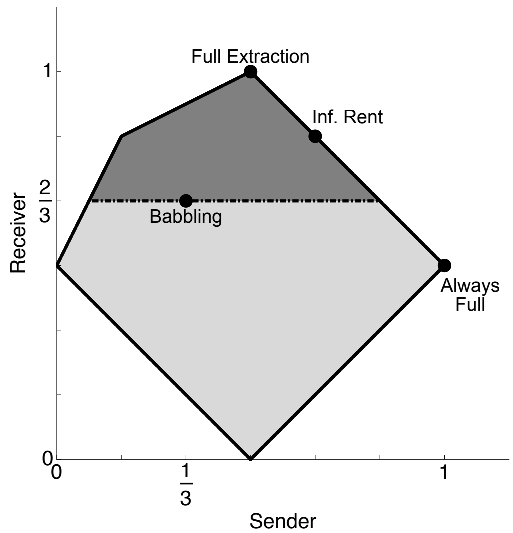

FIGURE 2. Feasible and IR Discounted-Average Payoffs

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

Strategy Definition 1 (Babbling)**.** *The babbling strategy is any strategy where: (i) the sender's message strategy in every period reveals no information about the state; and (ii) the receiver selects the uninformed best response (Partial) with certainty.*

A perfect Bayesian equilibrium (PBE) for the stage game requires us to specify a sequentially rational probability $\beta(a\,|\mathcal{H})$ for each available action a at each information set $\mathcal{H}$, and a system of beliefs $\mu(\theta\,|\mathcal{H})$ that is consistent for any information set reached with positive probability.

A full specification of a particular babbling PBE would therefore be given by a sender choice $\beta(\text{Invest}\,|\theta) = 1$ for both states, and a receiver choice of $\beta(\text{Partial}\,|\text{Invest}) = 1$ and $\beta(\text{None}\,|\text{Invest}) = 1$, with $\mu(\text{Good}\,|\text{Invest}) = \frac{1}{2}$ and $\mu(\text{Good}\,|\text{Don't Invest}) = 0$. This particular construction has pooling on the *Invest* message, leading to a sequentially rational response of *Partial* investment on the path; supported by off-path play of *None* in reaction to the *Don't Invest* signal. While there are many different PBE constructions for this game, where some can reveal information about the state, all PBE share the feature that *Partial* investment is chosen with certainty in both states.

Result 1. *Babbling is a PBE of the* Strangers *game.*

2.3. **Partners Game.** With fixed partners within the supergame, more-informative outcomes than the stage-game PBE are supportable as equilibria, so long as δ is large enough. These more-informative outcomes are possible in our *Partners* treatment through history-dependent play.

Figure 2 shows the set of feasible and individually rational expected payoffs for our stage game. The set of feasible payoffs is indicated by the entire shaded region, while the darker gray region is the set of individually rational payoffs. The efficient frontier of the figure runs from the extreme point $(1/2, 1)$, labeled *Full Extraction,* to the extreme point $(1, 1/2)$, labeled *Always Invest*. A particular outcome lies on the efficient frontier in our game if and only if the outcome has the receiver fully invests with certainty when the state is good. Investment behavior in the bad state defines the locus of points that trace out the frontier, where our constant-sum restriction $(u_0 + v_0 = 1)$ leads to the one-for-one downward gradient.

At the leftmost extreme point on the frontier we have full extraction by the receiver. That is, this is the payoff that would result if a fully informed receiver maximized their own payoff, choosing full investment in the good state and no investment in the bad state. At the other extreme of the frontier is the sender's most-preferred outcome, which is produced by the receiver fully investing with certainty in both the good and bad state, hence the label *Always Invest.*

Away from the frontier, the figure illustrates the payoffs from the stage game's babbling outcome, at the point $(1/3, 2/3)$. Babbling, which gives the receiver their individually rational payoff, remains a PBE of the *Partners* game. Of note in the figure is that while the receiver's individually rational payoff is equal to the babbling payoff, the sender's individually rational payoff is lower. This is because the receiver can enforce a payoff of zero on the sender (her minimum) by choosing no investment, without any requirement on information revelation. As such, the worst-case individually rational payoff pair for both players is given by the point $(1/12, 2/3)$.[15]

With fixed pairs, history-dependent strategies can be used to support more-efficient outcomes so long as δ is large enough, and standard folk-theorem constructions can be used to support continuation values from the set of individually rational payoffs. In fact, given a public randomization device, all payoffs in the individually rational set are attainable as equilibria for $\delta = 3/4$.[16] Restricting to simple pure-strategy outcomes, an inefficient, stationary PBE exists—the babbling equilibrium. However, in addition to babbling, we now focus on characterizing two forms of history-dependent equilibria for $\delta = 3/4$: i) Equilibria with full revelation by the sender and full extraction by the receiver; and ii) Equilibria with full revelation but without full extraction.

Full revelation-Full extraction. Full-revelation requires the sender to reveal the state at every information set along the path of play (appealing to natural language, by sending the message *Invest* in the *Good* state and *Don't Invest* in the *Bad* state). Full-extraction requires that the receiver maximizes given their consistent beliefs (choosing *Full* investment on *Invest*, and *None* on *Don't*

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

[16]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.

Invest). As the restriction fully pins down the path we turn to how this outcome can be supported off the path. One natural candidate to supporting full extraction is reversion to babbling.

Strategy Definition 2 (Full extraction with babbling reversion)**.** *The strategy has two phases:*

> **Full extraction:** *The sender fully reveals the state, the receiver fully extracts.*
> **Punishment:** *Both players follow the babbling strategy.*

The game starts in the Full extraction phase and remains there as long as there are no deviations. Upon any deviation, the game moves to the punishment phase.

However, for the parameters of the game, the *Full extraction-babbling* strategy is not part of a PBE of the *Partners* game, as the babbling outcome is not a harsh enough punishment to incentivize senders to reveal the state.

Result 2. *For $\delta = \frac{3}{4}$, the Full extraction-babbling strategy is not a PBE of the* Partners *game.*

If both agents follow the *Full extraction-babbling* strategy, the discounted-average payoff on the path of play is $({}^1\!/_2, 1)$. For the strategy to satisfy the one-shot deviation principle it must be that a sender observing $\theta =$*Bad* does not want to deviate. Given that a deviation yields a static benefit of 1 in period t instead of 0, for the deviation to be unprofitable the sender's continuation value $U^\star$ under the punishment must satisfy:

$$(1) \qquad (1-\delta)\cdot 0 + \delta \cdot \tfrac{1}{2} \geq (1-\delta)\cdot 1 + \delta U^\star.$$

So for $\delta = {}^3\!/_4$ we require that $U^\star \leq {}^1\!/_6$. The babbling equilibrium is clearly too high, as $U_B = {}^1\!/_2$, and so a harsher punishment is required to support full extraction.

The fact that reversion to babbling cannot be used to support full extraction ends up being a general feature of these games under a fairly standard concavity assumption for senders.[17]

Proposition. *Consider a sender-receiver game with a continuous state space $\Theta = [0,1]$, with the corresponding message and action spaces $\mathcal{M} = \mathcal{A} = [0,1]$. If (i) the sender has a concave state-independent preference $u(a)$ increasing in a, and (ii) the receiver has a payoff $v(\theta, a)$ uniquely maximized at $a = \theta$ under full information with an uninformed best response of $a_\emptyset \geq \mathbb{E}\theta$; then full extraction can not be supported by babbling reversion for any $\delta \in (0,1)$.*

[17]The proposition shows that full extraction can not be supported by babbling reversion when the sender's payoff is concave. In our specific parametrization—and setting Partial investment to be the exact intermediate between full and no investment—the sender's payoff is instead *convex* over the action, ${}^1\!/_2 \cdot u_S(\text{Full}) + {}^1\!/_2 \cdot u_S(\text{None}) > u_S(\text{Partial})$. However, the result holds in our parametrization at $\delta = \frac{3}{4}$. We note that were the sender's payoff more convex full extraction *could* be supported with reversion to a PBE. For example, if $u_0 = 0$, the babbling payoffs for the sender are low enough that full extraction can be supported by babbling reversion. Such a parametrization would provide the receiver with a simple, powerful punishment to support his first-best outcome. However, it would not allow for any information on the extent to which harsher punishments than babbling are selected.

Proof. Full extraction (FE) requires $a = \theta$ along the path, so the sender's expected continuation payoff on the path is $u_{FE} = \mathbb{E}u(\theta)$. In the babbling outcome the receiver chooses $a_\emptyset \geq \mathbb{E}\theta$ with certainty, and the sender's payoff is u_B. In any state $\theta < 1$ there exists a strict gain from a deviation (signaling some $\theta' > \theta$, possible by full revelation, and increasing the current period outcome). But alongside the strict contemporaneous gain from the deviation the change in continuation provides a weak benefit ($u_{FE} = \mathbb{E}u(\theta) \leq u(\mathbb{E}\theta) \leq u_B$ by Jensen's inequality and $u(\cdot)$ increasing). This result is independent of δ for all values where the continuation values are well-defined. □

Reversion to babbling cannot be used to support full extraction in typically modeled environments (quadratic loss for the receiver; concave preference for the sender). In many related environments (and our stage game) this same result will instead hold whenever δ is small enough.[18]

This is not to say that harsher punishments than babbling cannot sustain fully extractive outcomes. For a harsher punishment, the receiver could choose zero investment for a number of periods, thus pushing the sender's payoff towards zero. In order for this to happen the receiver needs to take actions that give him a temporarily lower payoff than his individually rational payoff (in our game the *None* action, which has an ex ante expectation of 1/2 to the receiver). In order for this action to be rational for the receiver, it must be that after a number of periods of this punishment, the game must eventually return to a point that gives the receiver more than his babbling payoff.

Corollary. *Full extraction is only supportable as a PBE with strategies that punish the sender more than the babbling outcome. Such punishments require receivers to at least temporarily take outcomes which are myopically sub-optimal, and where the punishment path must at some point return to greater revelation than babbling.*

The above points out the necessary components of punishments capable of supporting full extraction when babbling reversion does not. In terms of Figure 2, an individually rational punishment payoff to the left of the babbling outcome must be selected. But all such points are generated by combinations of play where the receiver has a myopic loss and where more-information than babbling is revealed.

While fully extractive outcomes *are* possible in equilibrium using pure-strategies, there are two reasons to be concerned about the predictive power of such constructions. First, such outcomes rely upon coordination on sophisticated punishments. Second, typical pure-strategy constructions will rely on the receiver punishing across a number of periods (for instance, through the *None* action) reducing his period payoff below the individually rational one. For such punishments to be incentive compatible for the receiver, an eventual return to revelation is required. Supporting full

[18]In games where the sender has a state-dependent payoff with a convex loss function with a bias term $b > 0$, there will typically be an upper bound $\delta(b) < 1$ for the result, where for $\delta > \delta(b)$ the game will admit fully extractive outcomes supported by babbling reversion.

extraction therefore necessitates the "forgiveness" of senders who have acted dishonestly and who have been harshly punished.

While the above speaks to the behavioral forces arrayed against the selection of full extraction, it **is** supportable as an equilibrium in our game.

Result 3. *For $\delta = \frac{3}{4}$, the Fully Extractive path is supportable with harsher punishments than babbling.*

This result can be demonstrated by construction. For our parametrization there are punishments that support full extraction where the receiver chooses K_1 rounds of None, followed by K_2 rounds of Partial investment followed by a return to the fully extractive path.[19]

Full revelation-Information rent. Above we showed that for the receiver to extract all of the informational gains generated by full revelation, the punishments must be more-severe and more-complex than reversion to babbling. However, fully informative equilibria do exist where revelation is supported by a babbling trigger. In these efficient PBE the receiver shares the generated surplus with the sender, despite being fully informed in the bad state. As such, the payoff path lies to the right of the fully extractive point on the efficient frontier of Figure 2. We will refer to the distance to the right of full-extraction as the "information rent" being given to the sender.

While there are many possible efficient information rent equilibria, we focus here on the below construction, a simple pure-strategy PBE.

Strategy Definition 3 (Information rent with babbling-trigger.)**.** *The strategy has two phases:*

> **Information-Rent:** *The sender fully revels the state, the receiver selects "Full" when $\theta =$Good and "Partial" when $\theta =$Bad.*
> **Punishment:** *Babbling strategy.*

The game starts in the information-rent phase and remains there as long as there are no deviations. Upon any deviation, the game moves to the punishment phase, where it remains.

The expected on-path payoffs from the *Information rent-babbling* strategy are illustrated in Figure 2 at the point $(2/3, 5/6)$, labeled as the *Information rent* outcome. The strategy provides the sender with an expected rent of $1/6$ per period, transferred as an additional payment of $1/3$ in each bad-state realization, relative to the receiver's myopic sequentially rational response. The outcome doubles the sender's payoff relative to babbling, while increasing the receiver's payoff by 25 percent.[20]

[19]The (K_1, K_2) punishment strategy is a PBE of the *Partners* game for $(K_1, K_2) \in \{(2,4),(2,5),(2,6),(3,2),(3,3)\}$
[20]Moreover, the receiver's revenue stream (a 1 or $2/3$ each period) stochastically dominates the stream under babbling ($2/3$ each period).

For the strategy to be supported as an equilibrium we need to construct punishment phase payoffs $(U^{\star},V^{\star})$ on any observed deviation. To satisfy the one-shot deviation principle it is necessary that neither the sender or receiver wish to deviate, which leads to the following requirements on the punishment continuation payoffs $(U^{\star},V^{\star})$:

$$(1-\delta)\cdot\tfrac{1}{3}+\delta\cdot\tfrac{2}{3}\geq(1-\delta)\cdot 1+\delta\cdot U^{\star} \qquad \Rightarrow U^{\star}\leq\tfrac{4}{9},$$
$$(1-\delta)\cdot\tfrac{2}{3}+\delta\cdot\tfrac{5}{6}\geq(1-\delta)\cdot 1+\delta\cdot V^{\star} \qquad \Rightarrow V^{\star}\leq\tfrac{13}{18}$$

This simple outcome path, where the receiver pays the information rent by selecting partial investment when the state is revealed to be bad, can therefore be supported as equilibrium of the repeated game by reversion to babbling on any deviation as $(1/3,2/3)\ll(4/9,13/18)$.

Result 4. *For $\delta=\frac{3}{4}$, the Information rent with a babbling trigger is a PBE of the* Partners *game.*

While the particular information rents equilibrium we describe is an example of a simple pure-strategy path supported by a babbling-trigger, other efficient outcomes with information rents for the sender—points along the efficient frontier in Figure 2—are also possible. For instance, an informed receiver could choose some fixed sequence of partial and no investment each time the bad state is signaled.[21]

2.4. **Hypotheses.** Comparing the predictions for the *Partners* and *Strangers* treatments, we formulate the following hypotheses.

Hypothesis 1. *Information transmission in the* Partners *game should be equal to or higher than in the* Strangers *game. Outcomes in the* Partners *game should be at least as efficient as in the* Strangers *game.*

In addition, the theoretical predictions indicate more-precise restrictions on behavior conditional on revelation.

Hypothesis 2. *For the* Partners *game, if senders reveal information, one of the following holds:*

(1) *Receivers fully extract, but punish dishonesty with investment choices of* None.
(2) *Dishonesty is punished by babbling reversion, but receivers provide on-path information rents to senders.*

Contrarily, for the *Strangers* treatment we expect the following:

[21]There are no equilibria at $\delta=3/4$ supportable with babbling reversion where the receiver chooses full investment when the bad state has been perfectly revealed. *Full* investment can only be chosen in equilibrium in the bad-state if the sender and receiver coordinate on less information provision.

Hypothesis 3. *For the* Strangers *game, information rents are not paid by receivers when bad states are revealed, and senders do not reveal information.*

The above hypotheses are derived under the assumption that senders do not face a disutility from lying. However, full revelation may be observed in the *Strangers* game if senders have an aversion to lying (Kartik, 2009). To capture this alternative in the model we can modify the preferences of the senders to be

$$\tilde{u}_S(\theta, m, a) = u_S(a) - \eta \cdot \mathbf{1}\{m \neq m_\theta\}$$

where $\eta > 0$ measures the degree of the lying aversion.

If lying aversion is high enough ($\eta > 1$), the fully informative outcome is a PBE. In other words, observing full revelation in the *Strangers* treatment is consistent with a large enough aversion to lying. However, lying aversion at this level would lead to full extraction being the equilibrium prediction for both the *Partners* and *Strangers* treatments. As a relative difference between the treatments our above hypothesis would remain true under lying aversion.

3. Strangers and Partners Treatments

For both the *Strangers* and *Partners* treatments we recruited a total of 46 subjects across three separate sessions (16, 16 and 14), and each subject participated only once.[22] A session consists of 20 repetitions of the supergame, where each repetition involves an unknown number of periods ($\delta = 3/4$).[23] Half of the subjects are randomly assigned to be senders and half to be receivers, where these roles are kept fixed throughout the session. In the *Strangers* (*Partners*) treatment subjects were randomly matched to a different participant at the beginning of each new period (supergame). Stage-game payoffs in the laboratory are multiplied by 3, so that the bad state involves a zero-sum game over \$3, and in the good state the efficient outcome allocates \$6 in total.[24] In what follows we will present our main results from these two treatments.[25]

Table 1 provides a summary of average behavior within our experimental stage games. The table provides empirical measures for the components of the stage-game PBE: (i) $\hat{\beta}(x\,|\mathcal{H})$, the average

[22]All sessions were conducted at the University of Pittsburgh and we used zTree (Fischbacher, 2007) to design our interface.

[23]We randomized the length of supergames and the draws of the states prior to conducting the sessions so that the comparison between the two treatments are balanced. Specifically, the randomization of session i of the *Strangers* treatment is identical to the randomization of session i of the *Partners* treatment ($i = 1, 2, 3$). In addition, senders in each session are matched across treatments to face an identical sequence of draws for the state θ across the entire supergame.

[24]Each 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.

[25]A coding mistake led to an error in the subject matching in periods two onward in the last supergame of the first *Partners* session. Even though subjects were unaware of this, we drop the relevant data for this supergame.

TABLE 1. Behavior and Empirically Consistent Beliefs: *Strangers* and *Partners*

Category	Empirical Frequency	Strangers		Partners			
		$t=1$	All t	$t=1$		All t	
Sender Behavior	$\hat{\beta}(m = \text{Invest}\,\vert\theta = \text{Good})$	0.992	0.972	0.996		0.976	
		(0.006)	(0.009)	(0.004)		(0.007)	
	$\hat{\beta}(m = \text{Don't}\,\vert\theta = \text{Bad})$	0.259	0.197	0.422	***	0.316	***
		(0.029)	(0.019)	(0.032)		(0.022)	
Consistent Belief	$\hat{\mu}(\theta = \text{Good}\,\vert m = \text{Invest})$	0.574	0.551	0.633	***	0.588	***
		(0.010)	(0.004)	(0.013)		(0.008)	
	$\hat{\mu}(\theta = \text{Bad}\,\vert m = \text{Don't})$	0.968	0.922	0.990		0.929	
		(0.022)	(0.026)	(0.010)		(0.021)	
Receiver Behavior	$\hat{\beta}(a = \text{Full}\,\vert m = \text{Invest})$	0.360	0.272	0.412	*	0.310	
		(0.024)	(0.017)	(0.026)		(0.019)	
	$\hat{\beta}(a = \text{Partial}\,\vert m = \text{Invest})$	0.605	0.685	0.569	*	0.635	
		(0.024)	(0.019)	(0.026)		(0.021)	
	$\hat{\beta}(a = \text{None}\,\vert m = \text{Don't})$	0.567	0.544	0.673		0.600	
		(0.066)	(0.039)	(0.048)		(0.040)	
	$\hat{\beta}(a = \text{Partial}\,\vert m = \text{Don't})$	0.383	0.394	0.276		0.346	
		(0.064)	(0.039)	(0.045)		(0.041)	

Note: Standard-errors (in parentheses) are derived from a bootstrap (of size 5,000) with supergame-level resampling across the 460 supergames for each treatment. Significance stars represent t-tests for equality of coefficients between the relevant *Partners* and *Strangers* entry against one-side hypotheses justified by theory (information-rents strategy vs. babbling): *** -99 percent confidence; ** -95 percent confidence; * -90 percent confidence.

rate at which subjects choose response x at information set $\mathcal{H}$; and (ii) $\hat{\mu}(\theta\,\vert m)$, the empirically consistent belief for a receiver, computed as the relative frequency of each state θ conditional on message m. In total our data contains 460 supergames for each treatment. The table breaks the choices in these supergames out by the initial response (the first period of each supergame, $t = 1$) and overall across all supergame periods.

The first set of results in Table 1 outline the sender's strategy, where we provide the proportion of honest (using natural language) responses after each state is realized.[26] The first row indicates that senders choose the *Invest* message in the good state at high rates. This is true in both the *Strangers* and *Partners* treatments, particularly for the initial response where the *Invest* message is selected in over 99 percent of supergames, where the overall response proportion is in excess of 97 percent in both treatments.

[26]The 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.

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 message *Don't Invest*, however, our results indicate that this is not the modal message selection. In the first period of *Strangers* supergames senders select the *Don't* message in reaction to a bad state 26 percent of the time. This declines to 20 percent when we look at all periods. For the *Partners* treatment the rate of honest revelation of the bad state does increase. However, at 42 percent in the first period of each supergame, the rate of honest revelation 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.[27,28]

Though honesty is not the norm in the bad state, the message *Invest* does convey some information. The next set of results in Table 1 makes precise how much information. Each entry $\hat{\mu}$ is the empirical likelihood for the claimed state conditional on each message. The measure is therefore the belief that would be 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 indicates a 63.3 percent probability of the good state in the first period of the *Partners* treatment. 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.

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 period of *Partners*' supergames). The expected period payoff to the receiver choosing no investment after a *Don't* message is therefore \$2.97 in *Partners*. The high 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 group of results in Table 1 present the last piece of the puzzle, how receivers in our experiments actually behave. 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.

[27]In the first period 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 period of the session.

[28]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.

Receiver's behavior in response to signals for the good state directly affect the resulting efficiency. Choosing full investment in the good state is necessary for an efficient outcome, but involves trading off a certain payoff $2 against a lottery over $3 and $0, with probabilities dependent on the belief that the sender is truthfully revealing the state.

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 period of a *Partners* supergame and 36 percent of decisions in the first period of *Strangers* supergames.

Though a majority of receivers respond to the *Don't* message with no investment, a large minority does choose the partial investment choice. However, examining the precise figures in Table 1, the rate at which partial investment is chosen is higher in the *Strangers* treatment. That is, from the aggregate data, subjects are more likely to make choices corresponding to the payment of an information rent in the one-shot interactions than they are in the repeated relationship.[29]

The figures in the table provide evidence for the following two findings:

Finding 1. *Aggregate behavior in the* Strangers *treatment mirrors the babbling outcome. Though there is some honest response by a minority of senders, outcomes are not well described by the lying-aversion prediction.*

Finding 2. *Aggregate behavior in our* Partners *treatment does show more revelation than the* Strangers *treatment; however, the difference is quantitatively small and the overall response is still best characterized as babbling.*

To summarize, behavior in both treatments is better represented as babbling. The possibility generated by history dependent strategies in the *Partners* treatment appears to have only a minor effect. In *Online Appendices B and C* we show that indeed there are significant differences in behavior across these two treatments in the expected direction. In particular, we do find evidence for history dependent strategies in the *Partners* treatment, that is absent in the *Strangers* treatment. However, though the differences between the *Partners* and *Strangers* treatments are statistically significant—where the directions mirror the theoretical predictions in Hypothesis 1—the size of the effects are small.

[29] The one-tailed tests for significance in Table 1 generated by the theory go in the opposite direction from the effect. In fact, if we were to state the opposite alternative we would be able to reject equivalence between the two treatments for receivers choosing *Partial* in response to *Don't* at 10 percent significance. From another point of view, the information at the aggregate level may be misleading with respect to subjects' strategic choices. In *Online Appendix C* we present a detailed analysis of strategies at the individual level.

The previous literatures on communication and repeated play in PD games suggested that outcomes in the *Partners* treatment could be highly efficient. Relative to these priors, and the economic sizes of the effects that we find, we characterize our findings as a null result here. Our next treatment examines the extent to which this null finding is driven by coordination failure.

4. Pre-Play Communication

One reason for the failure to find extensive efficiency gains in our *Partners* treatment is that tacit coordination in this environment is too hard. In our environment subjects have to coordinate on two levels. First, they need to coordinate on full revelation. If the sender does not reveal her information, efficiency gains are not possible. Second, they need to coordinate on the distributional issue that arises when information is revealed. While the first coordination problem (efficient vs. inefficient equilibria) is also present in a repeated PD, the second (selecting among efficient equilibria) is typically not. Moreover, in selecting among efficient equilibria two components of the strategy must be adjusted: a higher share for the receiver on-path needs to be accompanied by harsher punishments off-path. Difficulties in coordination over any of these margins may preclude subjects from achieving efficient outcomes. This raises the question: Would subjects in our *Partners* treatment be able to achieve better outcomes if provided with a device to relax their coordination problem?

To examine this question, our *Chat* treatment modifies the *Partners* repeated game to provide subjects with a powerful coordination device: pre-play communication. While identical to the *Partners* treatment for the first twelve supergames, in the last eight supergames of our *Chat* treatment sessions we allow subjects to freely exchange messages through a standard chat interface.[30] Recruiting exactly 16 subjects for each of our three *Chat* sessions, we use a perfect-stranger matching protocol in the last eight supergames with pre-play communication.[31] A chat interface is available to the matched participants for two minutes before the sender-receiver supergame begins, but not after the first period begins. As such, chat provides an opportunity for the subjects to coordinate on supergame strategies, but does not help as a channel for transmitting information on particular state realizations, or recoordination as the supergame evolves.

Alongside the same period-level data as our *Partners* treatment, *Chat* additionally provides chat-log data from subjects' pre-play conversations. To analyze these chat logs we designed a coding

[30] 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 eight supergames.

[31] Each subject interacts with participants in the other role from them for one chat supergame only. In addition, each new supergame match is guaranteed to have had no previous contact with others the subject previously engaged with. The design was chosen ex ante to allow us to isolate contagion effects for particular strategies, ex post, this ended up being moot.

TABLE 2. Behavior and Empirically Consistent Beliefs: *Partners* and *Chat*

Category	**Empirical Frequency**	**First 12 Supergames**			**Last 8 Supergames**		
		Partners		**Chat**	**Partners**		**Chat**
Sender Behavior	$\hat{\beta}\,(m = \text{Invest}\,\vert\,\theta = \text{Good})$	0.966	~	0.963	0.979	~	0.995
		(0.012)		(0.009)	(0.009)		(0.006)
	$\hat{\beta}\,(m = \text{Don't}\,\vert\,\theta = \text{Bad})$	0.365	~	0.361	0.264	***	0.813
		(0.029)		(0.024)	(0.033)		(0.029)
Consistent Belief	$\hat{\mu}\,(\theta = \text{Good}\,\vert\,m = \text{Invest})$	0.603	~	0.601	0.565	***	0.842
		(0.011)		(0.009)	(0.011)		(0.021)
	$\hat{\mu}\,(\theta = \text{Bad}\,\vert\,m = \text{Don't})$	0.914	~	0.907	0.948	~	0.994
		(0.026)		(0.022)	(0.032)		(0.008)
Receiver Behavior	$\hat{\beta}\,(a = \text{Full}\,\vert\,m = \text{Invest})$	0.364	~	0.413	0.222	***	0.830
		(0.024)		(0.024)	(0.025)		(0.039)
	$\hat{\beta}\,(a = \text{Partial}\,\vert\,m = \text{Invest})$	0.582	~	0.526	0.709	***	0.154
		(0.026)		(0.025)	(0.031)		(0.036)
	$\hat{\beta}\,(a = \text{None}\,\vert\,m = \text{Don't})$	0.623	~	0.638	0.531	***	0.225
		(0.041)		(0.042)	(0.071)		(0.035)
	$\hat{\beta}\,(a = \text{Partial}\,\vert\,m = \text{Don't})$	0.322	~	0.319	0.409	***	0.729
		(0.042)		(0.044)	(0.072)		(0.039)

Note: Standard-errors are in parentheses and are derived from a bootstrap (5,000 replications) with supergame-level resampling across supergames for each treatment-supergame-block. Significance stars represent t-tests for equality of coefficients between the relevant *Partners* and *Chat* treatments (two-sided as equilibrium sets are identical): ⋆ ⋆ ⋆-99 percent confidence; ⋆⋆ -95 percent confidence; ⋆ -90 percent confidence; ~Fail to reject at 10 percent significance or below.

protocol with 23 questions (available in the *Online Appendix A*). Two undergraduate research assistants (with no prior knowledge of our research hypotheses) independently coded each supergame chat log according to this protocol. Before analyzing the chat data, we examine the treatment effect from adding the coordination device, comparing the state/message/action data in *Partners* and *Chat*.

Table 2 provides aggregate statistics for the *Partners* and *Chat* treatments, broken out into supergames before the addition of pre-play communication in *Chat* (supergames 1–12) and after (supergames 13–20). A comparison of the two treatments for the first 12 supergames shows that the quantitative differences are small. In addition, our analysis of the data finds little evidence for information-rent strategies in either treatment over the first 12 supergames. We therefore focus on the differences in behavior in the last eight supergames.[32]

[32]In *Online Appendix C* we present an analysis at the individual level that reaches the same conclusion.

Though behavior in the first 12 supergames is similar, once we provide the coordination device, stark differences emerge between *Partners* and *Chat* (both statistically and economically significant) for the remaining eight supergames. Overall, *Chat* senders report the truth more than 80 percent of the time, where their honesty in the bad-state is increased by 50 percentage points over *Partners*. The effect of greater honesty in the senders' bad-state signals is that messages indicating the good state are more credible. This can be seen in Table 2 by examining the $\hat{\mu}\left(\theta = \text{Good}\,|m = \text{Invest}\right)$ row, which indicates that an *Invest* message is 30 percentage points more likely to correspond to a good state in *Chat* than *Partners.* Receivers' action choices in *Chat* reflect the greater information content given *Invest* messages, with full investment chosen 83 percent of the time in *Chat*, compared to just 22 percent in *Partners*.

The changes in behavior lead to large efficiency gains when pre-play communication is available. Conditional on a good state (the efficiency state), full investment is selected only 28.4 percent of the time in *Partners* compared to 86.1 percent in *Chat* (last eight supergames). However, alongside the increase in efficiency, we also observe a change in the distribution of payoffs when a bad state is revealed. In response to *Don't Invest* messages—highly correlated with the bad state in both treatments—we observe a 30 percentage point increase in the selection of partial investment in *Chat*. Given that the *Don't Invest* signal leads to subsequent feedback to receivers revealing the bad state in 99.4 percent of *Chat* periods, the increased selection of partial investment is not myopically sequentially rational. Instead, as we show below with the chat data, this choice is a direct effect of the large majority of subjects coordinating on the information-rents strategy.

Examining the coded chat data, and conservatively reporting averages only for data where the coders agree, we find that approximately three-quarters of complete chat exchanges have the sender and receiver explicitly discuss the path-of-play for the information-rents strategy. While punishments for deviations are only explicitly addressed in approximately one out of ten exchanges, when punishments are discussed they always refer to babbling reversion.[33] As an example of a chat interaction with a discussion of punishment consider the following exchange between a receiver subject (R_{180}) and a sender subject (S_{18}), where we have edited the exchange to match the economic labeling in the paper:[34]

> R_{180}: I'll trust you until you lie and then it's *[partial]* the whole way out.
> S_{18}: Hey want to work together on this?
> R_{180}: If you click *[Don't Invest]*, I'll go *[Partial]* so we both get something
> S_{18}: I will tell you all the honest computer decsions if you never click *[None]*

[33] *Online Appendix C* uses the Strategy Frequency Estimation Method (SFEM) of Dal Bó and Fréchette (2011) to provide estimates of the frequencies of strategies. A frequency of approximately 90 percent is attributed to the information-rents strategy once chat is introduced.

[34] The experiment had a left/middle/right (L/M/R) labeling for states, messages and actions. We replace that here with our economic labeling, indicated the edit with square brackets. All other text is verbatim.

S_{18}: instead when i mark *[don't]* click *[partial]*
S_{18}: deal?
R_{180}: no problem
R_{180}: deal

Though this chat is non-representative with respect to the receiver's articulation of the punishment, the discussion of the on-path part component is entirely representative.[35]

We note that the conversation quoted above, in which both subjects nearly simultaneously outline the information-rents strategy comes from the very first supergame where subjects had the chat interface. Examining the RA-coded data, we find that 50 percent of the first-instance chat exchanges mention the information-rents strategy, which grows to 86 percent by the final chat supergame.[36] While some subjects learn and adapt the information rent strategy from previous interactions, the fact that so many supergames discuss it at the very first opportunity indicates the extent to which subjects are aware of the strategy, but require pre-play communication to coordinate on it.[37]

In terms of efficiency, 91 percent of partnerships that discuss the information-rents strategy have a perfectly efficient supergame. In contrast, only 40 percent of chats that do not discuss the information rents strategy have fully efficient outcomes.[38]

Summarizing our findings from the chat treatment:

Finding 3. *While aggregate behavior in the* Partners *treatment mirrors the* Chat *treatment in the first 12 supergames, once chat is introduced large differences between the treatments emerge. The evidence from both observed behavior in* Chat *and the subject exchanges is consistent with successful coordination on the information-rents strategy driving increased efficiency.*

The fact that such a large proportion of subjects explicitly discuss the information rents strategy from the very first chat supergame, alongside the strong efficiency increase, strongly implicates coordination failure as the cause of low efficiency in *Partners*. Once subjects are provided with an explicit channel through which to coordinate they succeed, where without it they fail. The evidence from our *Chat* treatment indicates both that the complexity of the game is not overwhelming for subjects, and that the payoffs implemented provide enough incentive to spur coordination on more-efficient SPE.

[35] *Online Appendix D* contains all chat logs from our three *Chat* sessions.

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

[37] There are no systematic differences in who initiates the discussion of this strategy: 47 percent of supergames mentioning the information-rents strategy have the first mention by the sender, while 53 percent are initially driven by the receiver.

[38] For details, and evidence for statistically significant differences driven by the information rents strategy, see *Online Appendix A*, where we present a Tobit regression of the supergame efficiency on features of the pre-play discussion.

Below we introduce a further two treatments where we explore if it is possible to aid coordination without a direct channel for strategic collusion.[39] Removing the pre-play communication device, we instead provide two alternative coordination devices, maintaining the same effective incentives. Our first provides senders with a device to aid coordination on revealing the state. The second, provides a device to help receivers coordinate on distribution.

5. Strategic Coordination Without Chat: Two Alternatives

Reviewing the exchanges in *Chat* we examine two alternative coordination devices, testing two hypotheses regarding what the chat allows subjects to signal. The first hypothesis is related to difficulties in coordination over the honest revelation of information. In the absence of chat, senders might consider a bad-state realization in period one valuable with respect to coordination, as the sender can signal to the receiver her decision to coordinate on truthful revelation. In contrast, when the first-period realization is good, there is no way for an honest sender to separate from the *Always Say Invest* babbling strategy. In response to *Invest*, the receiver may directly select *Partial*. A vicious circle ensues if the sender increases their belief that the receiver is coordinated on *Always Partial*, choosing not to reveal bad states in subsequent periods. The first treatment we present in this section—our *Revelation* manipulation—allows senders to convey their intentions to reveal in every first period.

Our second hypothesis is that coordination difficulties are concentrated on distributional challenges. From this perspective, chatting is useful as it serves to reassure senders that they will receive rents for revealing the bad state. That is, chat serves to coordinate both parties on the strategy that *Partial* is chosen following *Don't*, despite sender and receiver knowing that the message reveals a bad state. Without the coordination channel, senders could expect to be taken advantage of if they reveal bad states, and so the *Invest* messages is sent in both states. Moreover, if receivers do choose *Partial* in response to a *Don't Invest* message, senders may be uncertain on whether the receiver is coordinating on information rents, or on a babbling outcome that chooses *Partial* investment in response to both messages. The second treatment we introduce—our *Distribution* manipulation—modifies the action space of the receiver to allow him to more clearly signal the provision of an information rent.

5.1. **Revelation Coordination.**

In this treatment, we change the timing of the *Partners* game to utilize a strategy method for senders. Before a state is selected, senders are asked to indicate a message choice for each possible state realization. After senders submit their state-dependent choices ($\mu_t : \{\text{Good}, \text{Bad}\} \rightarrow \{\text{Invest}, \text{Don't}\}$), the state θ_t is drawn for the period and the interface sends

[39] While 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). Our next two treatments remove the ability to offer promises.

the receiver the relevant message, $m_t = \mu_t(\theta_t)$. The receiver observes the message m_t realization as before and selects an action. A crucial difference is in the provided feedback: the receiver learns what message she would have received in the counterfactual state—that is the receiver's period feedback is effectively $(\theta_t, \mu_t(\text{Good}), \mu_t(\text{Bad}), a_t)$.

The manipulation is designed not to change the theoretical predictions in any substantive way, it simply allows senders to clearly signal their strategic intention to honesty reveal in the first period, regardless of the state's realization. Senders who decide to reveal (not reveal) can now be identified from the first period, where in *Partners* this happens only in bad states.

Table 3 presents the main aggregate behavioral responses for the treatment, where we provide data from the *Partners* treatment for comparison. We present the last eight supergames of each treatment to examine the longer-run effects. As is readily observable from the table, the *Revelation Device* treatment has little differences with the *Partners* treatment. Senders report the truth under the bad state approximately 30 percent of the time in the manipulation, slightly above the corresponding 26 percent for the *Partners* treatment. Moreover, the most-common receiver response to the messages *Invest* and *Don't* are, respectively, *Partial* and *None*. The majority of subjects in the treatment are again coordinated on the babbling equilibrium, with behavior quantitatively close to the *Partners* treatment across our measures.

The manipulation does produce a small efficiency gain: *Full Investment* is selected in the good state 35 percent of the time, relative to 22 percent in the *Partners* treatment, where this difference is significant at the 5 percent level. While allowing for the sender to reveal their strategy regardless of the realization of the state may lead to an increase in efficiency, the effect is relatively small. We summarize the finding next.

Result 5 (Revelation Device)**.** *Providing senders with a device to signal their coordination on full revelation at the end of every period does not lead to a substantial efficiency gain. Observed behavior is close to the* Partners *treatment.*

5.2. **Distribution Coordination.** According to the second hypothesis, part of the difficulty to achieving high efficiency in the *Partners* treatment is that subjects find it hard to coordinate on a commonly understood information rent being paid when the state is bad. In our *Distribution* manipulation, we modify the receiver's action space to make the possibility of sharing gains more visible and explicit. In each period the stage-game is identical to the *Partners* treatment except that after the receiver makes his action choice, he is asked to make a second choice prior to getting the

TABLE 3. Average Behavioral Choices and Consistent Beliefs: Revelation and Distribution treatments (last eight supergames)

Category	Empirical Frequency	Partners		Revelation			Distribution			
		$t=1$	All	$t=1$	All		$t=1$		All	
Sender Behavior	$\hat{\beta}\,(m = \text{Invest}\,\vert\theta = \text{Good})$	1.000	0.979	1.000	0.970		0.977		0.980	
		–	(0.009)	–	(0.007)		(0.015)		(0.009)	
	$\hat{\beta}\,(m = \text{Don't}\,\vert\theta = \text{Bad})$	0.404	0.264	0.429	0.312		0.689	$\star\star\star$	0.565	$\star\star\star$
		(0.050)	(0.033)	(0.037)	(0.033)		(0.047)		(0.040)	
Consistent Belief	$\hat{\mu}\,(\theta = \text{Good}\,\vert m = \text{Invest})$	0.627	0.565	0.647	0.587		0.757	$\star\star\star$	0.693	$\star\star\star$
		(0.020)	(0.011)	(0.021)	(0.013)		(0.028)		(0.019)	
	$\hat{\mu}\,(\theta = \text{Bad}\,\vert m = \text{Don't})$	1	0.948	1	0.902		0.967		0.966	
		–	(0.032)	–	(0.027)		(0.022)		(0.015)	
Receiver Behavior	$\hat{\beta}\,(a = \text{Full}\,\vert m = \text{Invest})$	0.375	0.222	0.367	0.304	$\star\star$	0.439		0.466	$\star\star\star$
		(0.041)	(0.025)	(0.041)	(0.027)		(0.051)		(0.039)	
	(with transfer)						0.097		0.054	
							(0.031)		(0.013)	
	$\hat{\beta}\,(a = \text{Partial}\,\vert m = \text{Invest})$	0.604	0.709	0.554	0.575	$\star\star\star$	0.500	$\star$	0.478	$\star\star\star$
		(0.041)	(0.031)	(0.041)	(0.034)		(0.051)		(0.040)	
	(with transfer)						0		0.009	
							–		(0.004)	
	$\hat{\beta}\,(a = \text{None}\,\vert m = \text{Don't})$	0.600	0.530	0.600	0.576		0.900	$\star\star\star$	0.872	$\star\star\star$
		(0.078)	(0.071)	(0.074)	(0.062)		(0.039)		(0.029)	
	(with transfer)						0.557		0.596	
							(0.063)		(0.048)	
	$\hat{\beta}\,(a = \text{Partial}\,\vert m = \text{Don't})$	0.350	0.408	0.356	0.378		0.100	$\star\star\star$	0.124	$\star\star\star$
		(0.076)	(0.072)	(0.072)	(0.057)		(0.039)		(0.028)	
	(with transfer)						0		0	
							–		–	

Note: Standard-errors are in parentheses and are derived from a bootstrap (5,000 replications) with supergame-level resampling across supergames for each treatment-supergame-block. Significance stars represent t-tests for equality of coefficients against the comparable Partners coefficient (two-sided as equilibrium sets are identical): $\star\star\star$-99 percent confidence; $\star\star$ -95 percent confidence; $\star$ -90 percent confidence.

period feedback. The second choice involves a transfer: the receiver can select to transfer $1 or $0 to the sender. Final payoffs are identical to the *Partners* treatment except for the transfer.[40,41]

The distribution manipulation allows participants to pay the information rents in another way. Along the information-rents path, the sender fully reveals the state, and the receiver selects *Full* with no transfer when the singled state is good. The difference from the original strategy is that when the receiver gets a bad state signal, they respond by selecting the myopic best response of *None* (instead of *Partial*). To pay the information rent the receiver instead transfers $1 to the sender. The final payoffs from these choices in the bad state are therefore $1 for the sender ($0 from the action, plus $1 from the transfer) and $2 for the receiver ($3 from their choice minus the $1 transfer). The transfer strategy therefore produces the exact same bad-state payoffs to both participants as when the receiver selects Partial investment and no transfer.

It should be noted that the addition of the transfer does expand the set of feasible payoff for the repeated game (see Figure A1 in the *Online Appendix* for a graphical depiction). However, the manipulation does not change the set of individually rational payoffs, as the receiver (sender) can still guarantee himself (herself) a payoff of $2 ($0). The main difference between the *Distribution* and *Partners* treatments is that the action space allows the receiver to pay the information rent explicitly, via a transfer. Under the modified information-rents strategy, the transfer is made in response to a *Don't* message signaling the bad state. The transfer of a dollar is made alongside a choice of no investment, showing both that the receiver followed the recommendation *and* wants to pay a rent.

The last two columns in Table 3 summarize aggregate level behavior in this treatment. A first observation is that senders are more likely to honestly reveal: the proportion of truthful revelation in the bad state increases by 30 percentage points relative to *Partners* (all periods). There is subsequent large effect on the empirically consistent beliefs, where in particular the invest message is approximately 20 percentage points more likely to reveal the good state. Finally, there is a large difference in receivers' choices. The proportion of choices leading to Full investment after an Invest message more than doubles to 46 percent.

Consistent with the information-rents equilibrium, there are very few transfers when the message is Invest. This is shown in the row below $\hat{\beta}\,(a = \text{Full}\,|m = \text{Invest})$, labeled "with transfer", where we provide the fraction of periods for which $\hat{\beta}\,(a = \text{Full}, x = \$1\,|m = \text{Invest})$. Overall, just five percent of choices in response to the Invest message make the $1 transfer alongside full investment, which is approximately one-in-nine of the Full investment choices. Transfers are almost never

[40]More formally, the action space for receivers is now $\hat{\mathcal{A}} = \{\text{Full, Partial, None}\} \times \{0, 1\}$, with a generic choice (a, x) being an action a and a transfer choice x, where the period payoffs (in dollars) to the sender and receiver are $3u(a) + x$ and $3v(a, \theta) - x$, respectively.

[41]Feedback at the end of the period also contains information on transfers.

made (< 0.01) alongside a Partial investment choice (after either message) or alongside None after the Invest message. However, contrasting with the very infrequent use of transfers elsewhere, we find that receivers *do* make transfers at substantial rates in the situation predicted by the modified information-rents strategy: in response to a *Don't Invest* message, alongside the no investment action.

Relative to the *Partners* treatment we observe a a large increase in the choice of *None* following the *Don't Invest* message in the *Distribution* treatment (a 34 percentage point increase to 87 percent. However the large majority of these zero-investment choices (68 percent) are coupled with a $1 transfer to the sender. Pooling together the two alternative information-rent responses (*Partial*/no transfer, and *None*/transfer) we find that 72 percent of responses in our manipulation pay an information rent in response to *Don't*. This represents a 29.8 percent point increase over *Partners* (significant at the 1 percent level). Overall, the aggregate statistics indicate that behavior is more-consistent with the information-rents equilibrium when the rent can be paid as an explicit transfer.

Examining the ensuing efficiency: Full investment is selected for 52.1 percent of good states. This figure represent more than double the efficiency in our *Partners* treatment, and approximately 60 percent of the efficiency level achieved in our chat treatment. The findings for this treatment suggest that a simple manipulation of the receiver's action-space that allows information rents to be explicit paid achieves substantial efficiency gains. Moreover, the observed behavior reflects the information-rents equilibrium supported by a babbling trigger. We summarize the finding for this treatment.

Result 6 (Explicit Rents)**.** *Modal behavior in the treatment where we provide subjects with a distributional coordination device is consistent with the information-rents equilibrium; the possibility of an explicit transfer greatly increases efficiency over our* Partners *treatment.*

6. CONCLUSION

Our paper examines the intersection of two important economic topics: the transfer of information, and the use of dynamic strategies. For the transfer of information an array of behavioral forces (lying aversion, bounded rationality) have been show to push outcomes in one-shot settings toward greater revelation than predicted by theory. For dynamic strategies, a large experimental literature on infinite-horizon PD games suggests that subjects readily utilize dynamic strategies to produce more-efficient outcomes. However, when we combine these two forces, a situation where honest response can be an efficient equilibrium outcome, we do not find behavior that is substantially more efficient than the babbling outcome.

Though we find evidence that subjects in our *Partners* treatments are capable of responding dynamically to more-cooperative choices of their partners, the overwhelming coordination is on the low-efficiency babbling outcome. Through additional treatments we show that coordination on efficiency can be very challenging. Subjects in the informed sender role are willing to truthfully reveal, but only when they are not punished for revealing information against their own myopic interest; that is when they are provided with an information rent. While it is possible for honesty to be compelled in equilibrium without an information rent through the (credible) threat of punishment for any dishonesty, such punishments are strategically complex. In contrast, in our experiments the only punishments we observe are much simpler ones: reversions to the stage game equilibrium.

Our paper, along with other work on dynamic games (see Vespa and Wilson, 2015) suggests an empirically motivated refinement for the analysis of environments with folk theorems. A restriction to history-independent equilibria—for example, the stage-game PBE in our setting, or the MPE in a dynamic game—seems to be too strong, as it will remove more-efficient equilibrium strategies that are openly articulated and used by subjects in our experiments and others. However, while we do find evidence for dynamic strategies that revert to history-independent equilibrium play on deviations, we find no evidence for subjects using harsher punishments than this.

Punishment strategies that force players down to their individually rational payoffs are analytically useful in proving the extent of what is achievable with a dynamic response, but they seem to have little predictive content. Instead, our paper motivates behaviorally restricted folk theorems that examine what can be achieved with reversion to the history-independent equilibrium. Rather than the individually rational payoff set, the prediction can be filtered down to the payoffs that can be supported using Nash reversion.

Importantly, while this restriction on punishments matches observations on behavior, the restriction's equilibrium implications on what types of outcomes can be sustained predicts our main finding: information rents are necessary for more-efficient outcomes.[42] This prediction based on simple punishments, is borne out both in observed on path behavior and via subjects' pre-play communication. Where subjects are not rewarded for their honesty in our experiments, they stop revealing information. To take a complementary line from *King Lear* to our opening quote "Nothing comes from nothing."

[42]Notice that we are not claiming that information rents are always necessary for efficiency to arise. It is certainly possible to find parametrizations where full extraction by the receiver can be supported with reversion to a stage-game Nash equilibrium (see footnote 17).

ECONOMETRICA SUBMISSION ATTACHMENTS

- Online Appendices A–C: Further Figures and Analyses
- Online Appendix D: Complete *Chat* treatment logs
- Online Appendix E: Representative Instrutions and Screenshots
- Disclosure Statements

References

Abeler, J., D. Nosenzo, and C. Raymond (2016): "Preferences for truth-telling," IZA Discussion Paper.

Abreu, D. and Y. Sannikov (2014): "An algorithm for two-player repeated games with perfect monitoring," *Theoretical Economics*, 9, 313–338.

Arechar, A., A. Dreber, D. Fudenberg, and D. Rand (2016): "The role of communication in noisy repeated games," Mimeo.

Baron, D. P. and D. Besanko (1984): "Regulation and information in a continuing relationship," *Information Economics and policy*, 1, 267–302.

Blume, A., D. V. DeJong, Y.-G. Kim, and G. B. Sprinkle (1998): "Experimental evidence on the evolution of meaning of messages in sender-receiver games," *American Economic Review*, 88, 1323–1340.

Brown, M., A. Falk, and E. Fehr (2004): "Relational contracts and the nature of market interactions," *Econometrica*, 72, 747–780.

Cai, H. and J. Wang (2006): "Overcommunication in strategic information transmission games," *Games and Economic Behavior*, 56, 7–36.

Camera, G. and M. Casari (2009): "Cooperation among strangers under the shadow of the future," *American Economic Review*, 99, 979–1005.

Charness, G. and M. Dufwenberg (2006): "Promises and partnership," *Econometrica*, 74, 1579–1601.

Cooper, D. J. and J. H. Kagel (2005): "Are two heads better than one? Team versus individual play in signaling games," *American Economic Review*, 95, 477–509.

Crawford, V. P. and J. Sobel (1982): "Strategic Information Transmission," *Econometrica*, 50, 1431–1451.

Dal Bó, P. and G. R. Fréchette (2011): "The evolution of cooperation in infinitely repeated games: Experimental evidence," *American Economic Review*, 101, 411–429.

Dal Bó, P. and G. R. Fréchette (2016): "On the Determinants of Cooperation in Infinitely Repeated Games: A Survey," .

Dickhaut, J. W., K. A. McCabe, and A. Mukherji (1995): "An experimental study of strategic information transmission," *Economic Theory*, 6, 389–403.

Dreber, A., D. Rand, D. Fudenberg, and M. Nowak (2008): "Winners don't punish," *Nature*, 452, 348–351.

Duffy, J. and J. Ochs (2009): "Cooperative behavior and the frequency of social interaction," *Games and Economic Behavior*, 66, 785–812.

Dutta, P. K. (1995): "A folk theorem for stochastic games," *Journal of Economic Theory*, 66, 1–32.

Ettinger, D. and P. Jehiel (2015): "An Experiment on Deception, Credibility and Trust," UCL working paper.

Fischbacher, U. (2007): "z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental Economics*, 10, 171–178.

Fischbacher, U. and F. Föllmi-Heusi (2013): "Lies in disguise—an experimental study on cheating," *Journal of the European Economic Association*, 11, 525–547.

Fudenberg, D., D. Levine, and E. Maskin (1994): "The folk theorem with imperfect public information," *Econometrica*, 997–1039.

Fudenberg, D. and E. Maskin (1986): "The folk theorem in repeated games with discounting or with incomplete information," *Econometrica*, 533–554.

Fudenberg, D., D. Rand, and A. Dreber (2010): "Slow to anger and fast to forgive: Cooperation in an uncertain world," *American Economic Review*.

Gneezy, U. (2005): "Deception: The Role of Consequences," *American Economic Review*, 95, 384–394.

GOLOSOV, M., V. SKRETA, A. TSYVINSKI, AND A. WILSON (2014): "Dynamic strategic information transmission," *Journal of Economic Theory*, 151, 304–341.

HURKENS, S. AND N. KARTIK (2009): "Would I lie to you? On social preferences and lying aversion," *Experimental Economics*, 12, 180–192.

KAGEL, J. H. AND P. MCGEE (2016): "Team versus Individual Play in Finitely Repeated Prisoner Dilemma Games," *American Economic Journal: Microeconomics*, 8, 253–76.

KARTIK, N. (2009): "Strategic communication with lying costs," *Review of Economic Studies*, 76, 1359–1395.

KATZWER, R. (2013): "rgsolve for Java," *mimeo*.

LAFFONT, J.-J. AND J. TIROLE (1988): "The Dynamics of Incentive Contracts," *Econometrica*, 56, 1153–75.

RENAULT, J., E. SOLAN, AND N. VIEILLE (2013): "Dynamic sender–receiver games," *Journal of Economic Theory*, 148, 502–534.

SOBEL, J. (1985): "A theory of credibility," *The Review of Economic Studies*, 52, 557.

SUTTER, M. (2009): "Deception through telling the truth?! experimental evidence from individuals and teams," *Economic Journal*, 119, 47–60.

VESPA, E. AND A. J. WILSON (2015): "Experimenting with Equilibrium Selection in Dynamic Games," Stanford working paper.

WANG, J., M. SPEZIO, AND C. 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, 984–1007.

WILSON, A. J. AND H. WU (2016): "At-will Relationships: How an Option to Walk Away Affects Cooperation and Efficiency," University of Pittsburgh working paper.