TI 2026-047/VII
Tinbergen Institute Discussion Paper

Leveraging private information when interests diverge: An experimental comparison of communication versus restricted delegation

Silvia Dominguez-Martinez[1]
Randolph Sloof[1]

1 University of Amsterdam and Tinbergen Institute

Tinbergen Institute is the graduate school and research institute in economics of Erasmus University Rotterdam, the University of Amsterdam and Vrije Universiteit Amsterdam.

Contact: discussionpapers@tinbergen.nl

More TI discussion papers can be downloaded at https://www.tinbergen.nl

Tinbergen Institute has two locations:

Tinbergen Institute Amsterdam
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 598 4580

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900

Leveraging private information when interests diverge: An experimental comparison of communication versus restricted delegation*

Silvia Dominguez-Martinez and Randolph Sloof

University of Amsterdam and the Tinbergen Institute

June 2026

Abstract

This paper investigates how managers choose between communication and delegation when informed workers may have different objectives. We study a laboratory principal-agent environment in which only the worker observes the payoff-relevant state of the world. In three equally likely states, preferences are aligned; in a fourth state, preferences are opposed. The ex ante degree of interest alignment is exogenously varied via the likelihood that the conflict state occurs. We allow managers either to retain authority, delegate authority, or restrict the worker's discretion under delegation. We find that the manager's choice between communication and delegation per se exhibits no significant monotone relationship with interest alignment, but conditional on delegation managers impose tighter restrictions when interests become less aligned. Under communication, workers almost always communicate truthfully when interests are aligned, but distort information when interests conflict; overall, messages become more informative and influential as ex ante interest alignment increases. The comparison with equilibrium benchmarks is nuanced: communication is more informative than the pure-strategy benchmarks predict, while the mixed-strategy equilibrium benchmark accounts for part of the observed aggregate pattern. The remaining communication patterns are qualitatively consistent with limited behavioral heterogeneity, especially a modest role for credulous managers, suggesting that small departures from the fully strategic benchmark may help sustain informative communication even when workers remain largely strategic. Overall, managers often retain authority and rely on communication; when they delegate, restrictions increase as interests become less aligned. Communication remains viable because workers distort information mainly when interests actually conflict.

JEL: C90, D80, M20.
Keywords: Delegation, Communication, Laboratory Experiment, Organizational Economics.

*Parts of this paper previously circulated as a working paper entitled "Communication Versus (Restricted) Delegation: An Experimental Comparison." We used Refine.ink and ChatGPT 5.5 to improve the line of argument and the clarity of our writing. The paper benefited greatly from the insightful comments received from Suzanne Bijkerk, Wooyoung Lim, Joel Sobel and Junze Sun. An earlier version has been presented at conferences and workshops in Bergen, Geneva, and Maastricht; we gratefully acknowledge the remarks from the respective audiences. We thank the priority research area Behavioral Economics of the University of Amsterdam for financial support. Corresponding author: S.DominguezMartinez@uva.nl.

1 Introduction

Apart from performance evaluation and reward systems, a key feature of organizational architecture is the assignment of decision rights. In deciding whether – and if so, to what extent – to delegate decision authority to lower-level employees, managers face a basic trade-off between a loss of control and a loss of information (cf. Jensen and Meckling (1992)). Lower-level employees often possess superior information about local conditions and are therefore better positioned to tailor decisions to the circumstances at hand. At the same time, their objectives may differ from those of the manager. Delegation then risks giving informed employees the discretion to choose actions the manager would not have chosen had she possessed the same information. Centralized decision-making avoids this direct loss of control, but because communication by a potentially biased worker is strategic, it typically entails a loss of information. The organizational problem is therefore how to make use of dispersed information when those who possess it may have different objectives: by retaining authority and relying on imperfect communication, or by delegating authority, possibly with restrictions on the worker's discretion.

This trade-off has received widespread theoretical attention. In the seminal cheap-talk setting of Crawford and Sobel (1982), divergent preferences limit the information that can be credibly transmitted under centralized decision-making. Within this framework Dessein (2002) shows that, when the agent's bias is commonly known and the incentive conflict is not too large relative to the manager's uncertainty, full delegation outperforms direct communication. The intuition is that the bias in information transmission under centralization is an order of magnitude larger than the bias in decision-making after delegation. Delegation is therefore predicted to be valuable even though it entails a loss of control. Moreover, delegation need not be all or nothing. If the manager can restrict the set of options available to the worker, she can preserve part of the worker's informational advantage while limiting the most costly consequences of discretion. In the standard theoretical framework, optimally restricted delegation is therefore the appropriate way to balance the use of local information against the need to retain control (see also Holmstrom (1984); Ottaviani (2000); Alonso and Matouschek (2008)).

The frequent observation of centralized decision-making suggests, however, that the theoretical case for delegation is not decisive. Indeed, Dessein's main result does not generalize to all relevant environments. In particular, if the worker's preference divergence from the manager depends on information privately observed by the worker, then the manager is uncertain not only about the relevant decision environment but also about the direction and severity of the conflict of interest. Communication may then become more informative, and the trade-off between centralization and (full) delegation can go either way (cf. Dessein (2002), Section 8.2; Ottaviani (2000); Rush, Smirnov, and

Wait (2010)).[1] Moreover, behavioral considerations may further improve the relative performance of communication. Experimental studies of signaling games, for example, typically find that cheap talk is far more informative than standard theory predicts; see e.g. Cai and Wang (2006), de Haan, Offerman, and Sloof (2015), and Blume, Lai, and Lim (2020) for a review. One plausible interpretation is that some senders are somewhat lying averse. Consistent with this idea, Kartik (2009, Proposition 4) shows theoretically that lying costs can improve information transmission and may make communication more attractive relative to delegation. A related possibility, and one that appears relevant in our setting, is that receiver behavior may also matter. If some managers are willing to take messages at face value, their presence may affect not only their own decisions but also the incentives of strategic workers. Even limited credulity may then be strategically amplified, making it worthwhile also for strategic workers to preserve a more informative language in situations where interests are sufficiently aligned, rather than distort communication more broadly. Finally, behavioral forces may also make restricted delegation less attractive than the standard benchmark suggests, for instance if workers dislike being constrained and are willing to punish the manager when they are, as in Falk and Kosfeld (2006). These considerations make the comparison between communication and restricted delegation an empirical question, and one that may depend on the degree of interest alignment.

We study this question in a laboratory experiment based on a simple principal–agent environment in which the worker is better informed than the manager and their alignment of interests depends on this private information (cf. Blume, DeJong, Kim, and Sprinkle (1998); Blume, DeJong, Kim, and Sprinkle (2001)). The manager chooses among three projects, but only the worker observes the payoff-relevant circumstances. In three states, preferences are perfectly aligned: manager and worker prefer the same project. In a fourth state, preferences are opposed: the manager prefers one project, the worker another, and the intermediate project provides a compromise. We vary the ex ante degree of interest alignment by changing the probability that this conflict state occurs. The manager can either retain authority and receive a cheap-talk message from the worker, or delegate the project choice to the worker. A key feature of our design is that delegation can be restricted: the manager may choose which subset of projects the worker is allowed to select from. The experiment therefore tests the comparative statics of interest alignment for the informativeness and influence of communication, the extent of restrictions under delegation, and the manager's choice between communication and delegation.

Our main experimental findings are as follows. Under delegation, managers respond to weaker in-

[1] Other key assumptions are that the sender is fully informed on the state and that the receiver has no private information of his own. Otherwise cheap talk may dominate delegation in some cases; see e.g. Agastya, Bag, and Chakraborty (2014) for a setting where the sender only observes one dimension of the (essentially two dimensional) state of the world and Garfagnini, Ottaviani, and Sørensen (2014) for a setup where the receiver has additional (private) information.

terest alignment by imposing tighter restrictions on the worker's discretion, in line with the theoretical comparative statics. Under communication, workers typically communicate truthfully when interests are aligned, while concealing (by distorting) information only when interests conflict. As a result, communication conveys more information than the relevant pure-strategy equilibrium benchmarks predict, especially when interests are weakly aligned. The communication patterns are qualitatively consistent with a role for limited behavioral heterogeneity. In particular, a small share of lying-averse workers and especially a modest share of credulous managers can generate 'strategic amplification': the presence of some managers who take messages at face value can make it worthwhile for strategic workers to preserve an informative language in situations where interests are aligned. Thus, the results do not suggest that communication works because workers are generally honest or managers generally naive. Rather, they point to a more limited and more subtle mechanism through which small departures from the standard benchmark can affect strategic communication.

Related literature. A large number of experimental studies on delegation already exist. One strand of the literature focuses on strategic delegation as a commitment device or as a way to shift responsibility.[2] Another strand focuses on carefully inferring individual's intrinsic valuation of decision authority and control beyond their instrumental benefits, thereby providing a rationale why in general people may be reluctant to delegate in the first place, see e.g. Bartling, Fehr, and Herz (2008) and Owens, Grossman, and Fackler (2014). Closer to our study are experiments that examine the informational role of delegation.[3] Fehr, Herz, and Wilkening (2013), Dominguez-Martinez, Sloof, and von Siemens (2014), and Sloof and von Siemens (2021) study settings based on Aghion and Tirole (1997), where the worker has to exert costly effort to become informed. In these environments, retaining authority or closely monitoring the worker weakens the worker's incentive to acquire information, leading to a loss of initiative that has to be traded off against the loss of control under delegation. In our experiment, by contrast, the worker is already informed. The central issue is therefore not how delegation affects information acquisition, but rather how the manager should make use of the worker's private information: through communication under retained authority, or through delegation, possibly with restrictions on the worker's discretion.

The closest paper to ours is Lai and Lim (2012), who study a communication-delegation setting in which the worker is fully informed and strategic communication under centralization leads to a loss of information.[4] They consider two games. In one, communication is much more attractive for the

[2]See e.g. Fershtman and Gneezy (2001), Hamman, Loewenstein, and Weber (2010), Coffman (2011), and Bartling and Fischbacher (2012).

[3]Other experiments consider settings with one head office and two divisions and focus on the tradeoff between improving coordination between the divisions through centralization versus facilitating adaptation to local circumstances via delegation; see e.g. Brandts and Cooper (2025), Evdokimov and Garfagnini (2019), and Hamman and Martínez-Carrasco (2023).

[4]Battaglini, Lai, Lim, and Wang (2019) experimentally test the informational theory of legislative committees by

manager than full delegation; in the other, the manager is better off delegating. Their experimental results largely support these predictions, although many managers still refrain from delegating in the game where delegation is predicted. Lai and Lim (2012) interpret this under-delegation in a level-k framework: managers expect less-than-fully strategic workers to reveal useful information.

Our experiment differs from Lai and Lim (2012) in four main respects. First, and most importantly, we allow for restricted delegation, so that the manager can use the worker's information while excluding actions that are particularly harmful to her. Second, whereas Lai and Lim (2012) compare two games that differ in payoff structure and action space, we vary the degree of interest alignment directly, holding the underlying payoff structure and choice set fixed. This allows us to test comparative statics for communication, delegation, and the manager's organizational choice. Third, we use a richer message space: workers can either make factual statements about the state or recommend a project. This connects to experimental evidence that the form of communication matters, for instance because subjects may use vague messages to avoid outright lies, advice may induce different levels of trust and trustworthiness than information sharing, and self-constructed messages may have stronger effects than selected messages with similar content (see, e.g.,Serra-Garcia, Van Damme, and Potters (2011), Özer, Subramanian, and Wang (2018), and Qin, Wang, and Wu (2024)). Finally, besides considering level-k reasoning as a possible explanation for informative communication, we include incentivized measures of workers' lying aversion and managers' credulity. These measures are not strong individual-level predictors of behavior under centralization. At the aggregate level, however, the communication patterns are qualitatively consistent with limited behavioral heterogeneity: already a modest share of credulous managers can help sustain informative communication when interests are aligned, while workers distort information when interests conflict. In this interpretation, credulity matters less as widespread managerial naivete than as a possible source of strategic amplification.

The remainder of this paper is organized as follows. Section 2 presents the theoretical model and its (comparative statics) predictions. In Section 3 we describe our experimental design while Section 4 presents our main results. In Section 5 we discuss how various behavioral forces may explain some of the observed deviations from theoretical (point) predictions. Section 6 concludes.

2 The model and its predictions

Below we first describe the strategic situation that we consider and then turn to the equilibrium predictions assuming expected payoff maximization. The final subsection summarizes the comparative statics predictions that we aim to test with our experimental hypotheses.

comparing the "closed rule" (essentially delegation) with the "open rule" (cheap-talk communication).

2.1 Strategic situation

We consider a game between a manager and a worker. The manager has the initial decision authority which project $a \in \{A, B, C\}$ to implement. The optimal project choice depends on the state of the world $t \in \{1, 2, 3, 4\}$. Ex ante states 1, 2, and 3 each occur with a probability of $\frac{1}{3}q$ while state 4 occurs with probability $1 - q$ (for $0 < q < 1$). The worker knows the actual state, the manager only its prior distribution. In the 'congruent' states 1, 2, and 3 preferences of manager and worker are perfectly aligned; in these states both prefer project A, B, and C, respectively. In the 'conflict' state 4 preferences are completely opposed. The manager then prefers project A while the worker prefers C, and the in between project B can be seen as a compromise solution. Parameter q can thus be interpreted as an ex-ante measure of interest alignment (cf. Pitchik and Schotter (1987)); the larger q is, the higher the interest alignment between manager and worker is.[5]

In our experimental implementation of the above strategic situation, payoffs over state-action pairs are given by Table 1 below:[6]

Table 1: Payoffs over state-project pairs

$Pr(t)$	t	A	B	C
$q/3$	1	$80, 80$	$0, 0$	$0, 0$
$q/3$	2	$0, 0$	$100, 100$	$0, 0$
$q/3$	3	$0, 0$	$10, 10$	$100, 100$
$1 - q$	4	$120, 10$	$80, 80$	$10, 120$

Notes: The first payoff belongs to the manager, the second payoff belongs to the worker

In the absence of any information, the manager would either choose project A or B, depending on whether q is low ($q < \frac{4}{5}$) or high ($q > \frac{4}{5}$). Project C is then always dominated by project B in expected payoff terms. Because the manager's preferred project depends on the state, he would like to take a more informed decision. One possibility to do so is full or partial delegation of the decision right to the worker. Another possibility is that the worker sends a cheap talk message about the state before the manager takes a decision ('centralization/communication'). Of these two options, the

[5]Pitchik and Schotter (1987) study a discrete model of strategic information transmission with two states and two actions. The preferred action of the receiver varies with the state, the preferred action of the sender does not (although the intensity of his preferences does). The state where preferences over actions coincide occurs with probability r. Pitchik and Schotter compare their setup with the framework of Crawford and Sobel and argue that r can be interpreted as a measure of (ex ante) interest alignment.

[6]The general structure of payoffs has been chosen such that: (i) under delegation the manager may either consider to exclude no project, to exclude project C, or to exclude both B and C, whereas (ii) under centralization a worker knowing that $t = 4$ always has an incentive to pretend being $t = 3$. The exact payoff levels are subsequently chosen to draw the predictions in the next subsection sufficiently far apart.

manager prefers the one that yields her the most in expected payoff terms.

2.2 Equilibrium predictions

To investigate the theoretical tradeoff the manager faces between delegation and centralization, below we first consider the delegation and centralization case separately. After that we study the manager's optimal choice between the two.

2.2.1 Delegation

Consider the delegation case and let $\Delta \subseteq \{A, B, C\}$ denote the set of projects the manager allows the worker to choose from. A selfish worker then chooses the project from Δ that yields him the most. Anticipating this, the manager's optimal level of delegation follows from a straightforward comparison of expected payoffs (proofs are relegated to Appendix A):

Proposition 1 *The optimal level of delegation for the manager equals:*[7]

(a) *If $q > \frac{7}{10}$, then Full Delegation (referred to as (Del-(3)): $\Delta = \{A, B, C\}$;*

(b) *If $\frac{12}{23} < q < \frac{7}{10}$, then Restricted Delegation (Del-(2)): $\Delta = \{A, B\}$;*

(c) *If $q < \frac{12}{23}$, then No Delegation (Del-(1)): $\Delta = \{A\}$.*

Proposition 1 is intuitive. If preferences are well aligned (q high), the manager is best off by delegating full decision power. In case preferences are badly aligned (q low), the manager should effectively not delegate at all and fully restrict the worker's "choice". When interest alignment is in between (q intermediate), restricting the worker to choosing either A or B is optimal. Overall, in case of delegation the worker is thus predicted to obtain more discretionary power the better interests are aligned.

2.2.2 Centralization

Next consider the centralization/communication case. Here the manager always takes the project decision, after having received a cheap talk message from the worker about the state of the world. Without loss of generality we assume there are four possible messages $m \in \{m_1, m_2, m_3, m_4\}$. For ease of reference, these messages can be interpreted as literal statements about the state. Alternatively, messages m_1, m_2, and m_3 could be interpreted as saying that the best project for both is A, B, or C

[7] Throughout Propositions 1 to 3 we state the cases for generic, non-boundary values of q. At cutoff values, one or more incentive constraints bind and the equilibrium outcomes from adjacent regions coexist. Thus, for example, at $q = 7/10$ both Del-(2) and Del-(3) are optimal, and analogous coexistence applies at the other cutoffs. These boundary cases are immaterial for the comparative statics predictions and for our experimental treatments, which all lie away from all the theoretical cutoffs.

respectively, and m_4 that there is a conflict of interests. As is standard in cheap talk games, messages only get their meaning in equilibrium though.

As is well-known, cheap talk games allow for a multiplicity of equilibria; see Sobel (2013) for an informative discussion. In deriving our theoretical predictions we therefore focus on the arguably more plausible equilibria by employing the neologism-proofness refinement of Farrell (1987).[8] In the present game, the refinement has two relevant implications. First, it rules out equilibria in which type $t = 1$ cannot secure that project A is chosen. Second, when multiple pure-strategy equilibria exist side by side, which happens for higher values of q, it rules out the less informative among them. As ex ante both worker and manager prefer the equilibrium in which the larger amount of influential information is transmitted, the equilibrium selection here can thus also be justified by assuming that, among the pure-strategy equilibria, parties coordinate on the most informative one. Proposition 2 below characterizes all neologism-proof equilibrium outcomes in our game (where for ease of reference we use the above labeling of messages; obviously these equilibrium outcomes could also be sustained by permuting the messages).

Proposition 2 *The set of neologism proof equilibrium outcomes in the communication game corresponds to:*

- **(a)** *For $q > \frac{33}{43}$, Communication (Com-(3)): The worker sends message m_1 when $t = 1$, m_2 when $t = 2$, and m_3 if $t = 3$ or $t = 4$. In response the manager chooses A, B, and C respectively;*
- **(b)** *For $\frac{12}{23} < q < \frac{33}{43}$, Limited Communication (Com-(2)): The worker sends message m_1 when $t = 1$, and m_2 if $t = 2$, $t = 3$ or $t = 4$. In response the manager chooses A and B, respectively;*
- **(c)** *For $q < \frac{12}{23}$, No Communication (Com-(1)): The worker always sends message m_1 irrespective of its type and the manager chooses A in response;*
- **(d)** *For $\frac{66}{141} < q < \frac{33}{43}$, Communication using mixed strategies (Com-(3-mixed)): The worker sends message m_1 when $t = 1$, m_2 when $t = 2$, m_3 if $t = 3$, and employs a mixed strategy when $t = 4$, sending m_2 with probability $\tau_2 = \frac{33-43q}{33(1-q)}$ and m_3 with probability $\tau_3 = \frac{10q}{33(1-q)}$. The manager chooses A after m_1, B after m_2, and after m_3 she chooses A with probability $\alpha_{3A} = \frac{4}{11}$ and C with probability $\alpha_{3C} = \frac{7}{11}$.*

Because the $t = 4$ type always has an incentive to pretend being a $t = 3$ type, full information revelation cannot occur in equilibrium. Focusing on pure strategy equilibria first, type $t = 3$ thus

[8] In the present context neologisms are (out-of-equilibrium) messages that are assumed to have a literal meaning of the form: "choose project a, because my type belongs to set S". A neologism is credible if: (i) all types in S prefer a over the project they receive in equilibrium, (ii) all types outside S prefer their equilibrium project over a and (iii) choosing a is a best response for the manager when restricting the support of his prior beliefs to S. An equilibrium is neologism proof if no credible neologisms exist.

always pools with $t = 4$. Whether the other types can separate themselves from this pooled set depends on the alignment of interests. If preferences are well aligned (q high) they can do so and all three projects are chosen in equilibrium, whereas if preferences are badly aligned (q low) they cannot and only project A is chosen. In case q is intermediate, effectively only $t = 1$ can separate itself from the rest.

Part (d) of Proposition 2 reveals that there also co-exists a mixed strategy equilibrium for a range of intermediate q values. In this equilibrium type $t = 4$ mixes between pretending to be $t = 2$ and getting compromise project B for sure, and pretending to be $t = 3$ after which the manager mixes between A and C. The former requires that q is not too low (for otherwise the manager would prefer A after m_2), while the latter requires q not to be too high (for otherwise the manager prefers C over mixing after m_3). The mixed equilibrium is more informative than the pure strategy equilibrium that exists at the same time (either *Com-(2)* or *Com-(1)*) and yields both the manager and the worker the most in expected payoff terms.[9]

Proposition 2 reveals that, independent of whether or not parties are able to coordinate on the mixed strategy equilibrium, communication is (weakly) more informative when interests become better aligned. As a result, the manager is more likely to be influenced by the worker's message when q increases. Overall, under centralization workers' messages are predicted to be both more informative and influential the higher the level of interest alignment is.[10]

2.2.3 Delegation versus centralization

Anticipating the equilibrium outcomes described in Propositions 1 and 2, a selfish manager is expected to choose the organizational structure that yields her the most. This comparison depends on whether manager and worker are able to coordinate on the mixed strategy equilibrium under centralization. Proposition 3 below summarizes this.

Proposition 3 *Comparing delegation with communication in terms of expected payoffs for the manager, it holds that:*

(a) *If* $q > \frac{33}{43}$*: Del-(3) = Com-(3);*

(b) *If* $\frac{7}{10} < q < \frac{33}{43}$*: Com-(3-mixed) > Del-(3) > Com-(2);*

[9] Note that messages induce posterior supports over the state space that the manager can distinguish in equilibrium. In *Com-(3-mixed)*, these supports are $\{\{1\}, \{2, 4\}, \{3, 4\}\}$, which form an overlapping cover rather than a partition, since state $t = 4$ can generate two different messages in equilibrium. In *Com-(2)*, the corresponding supports are coarser and form the partition $\{\{1\}, \{2, 3, 4\}\}$. In our empirical analysis we use as measure of informativeness the association between state and message; the higher this association (ranging between 0 and 1) is, the more informative the message of the worker is.

[10] In our empirical analysis we use as measure of influentiality the association between message and action (ranging between 0 and 1).

(c) *If* $\frac{12}{23} < q < \frac{7}{10}$*: Com-(3-mixed)* $>$ *Del-(2)* $=$ *Com-(2);*

(d) *If* $\frac{66}{141} < q < \frac{12}{23}$*: Com-(3-mixed)* $>$ *Del-(1)* $=$ *Com-(1);*

(e) *If* $q < \frac{66}{141}$*: Del-(1)* $=$ *Com-(1).*

Note that the three different delegation outcomes in Proposition 1 (and thus their expected payoffs) correspond to the respective pure communication equilibria in Proposition 2. From an equilibrium outcome perspective delegation and communication thus appear quite similar. There are two important differences, however, which make that they can lead to different outcomes. First, the range of interest alignment levels for which the respective pure equilibria exist do not fully coincide. In particular, the Com-(3) equilibrium puts a stronger requirement on the level of interest alignment ($q > \frac{33}{43}$) than the Del-(3) outcome does ($q > \frac{7}{10}$). Here the intuition is as follows. By fully delegating decision authority, the manager can de facto ex ante commit to an outcome where project A is chosen if $t = 1$, B if $t = 2$, and project C if either $t = 3$ or $t = 4$. The lower bound on q (equal to $\frac{7}{10}$) just reflects whether such a commitment is worthwhile to make. Under communication, however, the manager should have an explicit incentive to choose C (rather than A) when she knows that either $t = 3$ or $t = 4$. For this to be self-enforcing type $t = 3$ should be the more likely one of the two, yielding an additional and more tight requirement on q. This also explains why, if we focus on pure strategy equilibria only, delegation (weakly) outperforms communication. Delegation allows for commitment and thereby relaxes some potentially binding self-enforcement constraints.

As Sobel (2013, fn. 7) notes, most applications of sender-receiver games assume that the receiver has a unique best response to each distribution over types, ruling out mixed equilibrium strategies for the receiver by design. This also holds for the analysis in Dessein (2002) that builds on the standard Crawford-Sobel framework (Crawford & Sobel, 1982). Focusing on pure strategy communication equilibria only, Proposition 3 is reminiscent of Dessein's finding that delegation outperforms communication. Delegation has the advantage that (for intermediate values of q) decision making is more sensitive to the information available than strategic communication is.

In contrast to most applications, in our game the manager does not have a unique best response to every possible type distribution. Hence a second difference between delegation and centralization in our setup is that under the latter, for intermediate values of q a mixed equilibrium exists as well. This Com-(3-mixed) equilibrium yields the manager more than any form of delegation does. Here the underlying intuition is that the overall equilibrium relation between state and project choice becomes stochastic when the manager uses a mixed strategy, and this "strategic ambiguity" may benefit her in disciplining the worker's information provision. In contrast, a worker who is granted decision power will never employ a mixed strategy when choosing the project, given that he perfectly knows the state (and has a strict preference order over projects). Ruling out strategic ambiguity through delegation may then be restrictive when the optimal mechanism is stochastic (as holds in our setting).[11]

Per Proposition 3 the manager is predicted to be indifferent between delegation and centralization for high and low levels of interest alignment. And for intermediate levels the comparison can go either way. In particular, if under centralization manager and worker would be able to coordinate on the mixed strategy equilibrium, then centralization is predicted to occur. But if not, delegation is more likely. The overall equilibrium predictions thus do not provide a directed hypothesis on how the choice for a particular organizational structure varies with interest alignment. It is important to note, however, that Proposition 3 compares centralization with *optimal* delegation. The nature of delegation (if chosen) thus changes if interests become better aligned, i.e. there is a shift at the intensive margin by giving the worker more freedom of choice. For low levels of interest alignment the manager strictly prefers centralization over *full* (unrestricted) delegation.[12] Only if interests become sufficiently aligned, full delegation is predicted to be potentially chosen. Therefore, comparing centralization with full delegation like in Dessein (2002), does yield a clear prediction that a choice for the latter becomes more likely when q increases.

2.3 Hypotheses

The comparative statics predictions that follow from the equilibrium analysis in the previous subsection and that we aim to test with our experiment can be summarized as follows.

H1 In case of delegation, managers put less restrictions on the worker's choice set for higher levels of interest alignment.

H2a In case of centralization, worker's messages are (weakly) more informative for higher levels of interest alignment.

H2b In case of centralization, worker's messages influence manager's behavior (weakly) more for higher levels of interest alignment.

H3a The manager's choice between delegation and centralization per se does not vary systematically with the level of interest alignment.

H3b A choice for *full* delegation (instead of centralization) becomes more likely the higher the level of interest alignment.

[11] In a previous version of the paper we derived the optimal mechanism for our setting and show that it is stochastic for intermediate values of q.

[12] This follows because for $q < \frac{7}{10}$ the manager strictly prefers restricted delegation (either *Del(2)* or *Del-(1)*) over full delegation (*Del-(3)*) and centralization induces an equilibrium outcome that either corresponds to this or is better (viz. *Com-(3-mixed)*).

The above hypotheses are based on equilibrium predictions under expected payoff maximization and do not take other behavioral forces into account that may affect the tradeoff between delegation and centralization. For instance, workers' lying aversion and managers' credulity may lead to more informative communication under centralization than predicted by Proposition 2. In a similar vein, workers' (negative) reciprocity may lead to hidden costs of control in case the manager restricts the worker's choice set under delegation. Although such behavioral forces may clearly play a role and impact effect sizes, they are unlikely to overturn the general direction of the above predicted comparative statics effects. In our Discussion section we return to considering such behavioral forces when interpreting our experimental findings.

3 Experimental design

As the purpose of our experiment is to test the comparative statics of centralization and delegation with respect to the level of interest alignment, we considered three different q values: 'strong' interest alignment ($q = \frac{3}{4}$), 'medium' interest alignment ($q = \frac{3}{5}$), and 'weak' interest alignment ($q = \frac{2}{5}$). These belong to cases (b), (c) and (e) in Proposition 3, respectively. We chose these values to have all three optimal delegation levels of Proposition 1 represented, while at the same time having an instance where full delegation might potentially be strictly preferred over centralization (viz. case (b) in Proposition 3 when focusing on pure equilibria).

The data was collected in two waves.[13] The general setup and the essential features of the experiment were kept constant across waves. In particular, each session of both waves consisted of four incentivized parts and an ex post questionnaire. In parts I and II subjects first gained experience with centralization and delegation in isolation, before the manager actually had to choose between these in part III. In Part IV we measured lying aversion of the worker and credulity of the manager. Overall earnings equaled the sum of earnings in the four parts. Subjects kept the same role (either manager or worker) throughout the experiment. The experiment was framed in an organizational setting; see Appendix F for sample instructions. We used a conversion rate of 250 points for 1 euro. The experiment was conducted at the University of Amsterdam and was programmed using the z-tree programming package by Fischbacher (2007).

Below we discuss the specific details of wave 1 and after that point out for wave 2 (only) the aspects in which it differed.

[13]The first of these two waves was already conducted long time ago (in the summer of 2014) when pre-registration was uncommon; we did not pre-register our experiment at the time. One of the main objectives of pre-registration we feel is to clearly distinguish between confirmatory and more exploratory research. In this paper, we build on existing work in the literature to develop a simple formal model that guides the formulation of our hypotheses. The resulting analysis thus has a clear confirmatory structure, which arguably limits the incremental value of pre-registration.

Table 2: Experimental design

q	# ses.	# subj.	# mg	S / R	Part I	part II	Part IIIa	Part IIIb
Wave 1								
0.75	3	66	7	S & R	C	D	C vs. D	C vs. D
0.75	2	54	6	S & R	D	C	C vs. D	C vs. D
0.6	2	60	6	S & R	C	D	C vs. D	C vs. D
0.6	2	50	6	S & R	D	C	C vs. D	C vs. D
Wave 2								
0.6	3	68	7	S	D	C	C vs. D	C vs. full D
0.6	3	62	6	R	D	C	C vs. D	C vs. full D
0.4	2	52	6	S	D	C	C vs. D	C vs. full D
0.4	2	54	6	R	D	C	C vs. D	C vs. full D

Notes: In all sessions of both waves, lying aversion and credulity were measured in part IV. # ses. = number of sessions, # subj. = number of subjects, # mg = number of matching groups. S / R reflects whether messages under centralization could be formulated in terms of State or Recommendation. The analysis of C versus D (where under delegation restrictions can be imposed) reported in Section 4.3 mainly focuses on part IIIa data; therefore 'C vs. D' in part IIIb of wave 1 is put in lighter print. 'C vs. full D' refers to the choice between centralization and full delegation, where under the latter no restrictions can be imposed.

Wave 1

In the first wave we considered two different values of q, varied between subjects: $q = 0.75$ and $q = 0.6$. In total 230 subjects participated, divided over 9 different sessions. The number of participants per session ranged from 20 to 30. In each session we had either two or three matching groups of either 8, 10, or 12 subjects. Subjects were only matched within their matching group (with random rematching every period); matching group averages can thus be taken as strictly independent observations. Table 2 also lists the overall number of matching groups. Sessions lasted around two hours. Average earnings equaled 20 euros, with a minimum of 13.8 euros and a maximum of 28.3 euros.

Parts I and II of the experiment gave subjects the opportunity to gain experience with Centralization and Delegation in isolation. In the Centralization game the subject in the role of manager took the project decision, but before doing so the subject in the role of worker sent a message about the actual state (of which only the worker was informed). We restricted the worker to use one of the following eight messages: "The state is t" for $t = 1, 2, 3, 4$, "I recommend project P" for $P = A, B, C$ and "I make no recommendation". Subjects were explicitly informed in the instructions that: "The set of available messages does not depend on the actual state; so irrespective of the actual state, the worker can always choose one of the above eight messages." Column 'S / R' in Table 2 reflects that in wave 1 subjects could formulate messages both in terms of state and in terms of recommendation (as indicated by 'S & R'). In the Delegation game the worker took the project decision, but before doing so the manager could restrict the worker's choice options. To account for potential order effects, half of the sessions started with Centralization (C) in part I, the other half with Delegation (D) in part I. In part III the manager first had to choose between Centralization and Delegation and then the corresponding subgame was played. (In wave 1 parts IIIa and IIIb were identical; this distinction only matters for wave 2, see below.)

Each of the parts I, II and III consisted of 20 periods. In each period the manager and the worker were anonymously and randomly rematched within their matching group. At the end of each period, a summary of the manager's and worker's decisions and the resulting payoffs in that period was shown to them. These payoffs equaled the number of points of the implemented project given the true state of the world as reflected in Table 1. Subjects did not receive information on the behavior of other managers and workers. Furthermore, before making any decision in part III, subjects received an overview of the decisions made in the first two parts of the experiment. The overall payoffs of each subject in each part was equal to the sum of points earned in all 20 periods.

In part IV we used an adjusted version of the experiment by Gneezy, Rockenbach, and Serra-Garcia (2013) to measure lying aversion of workers as well as credulity of managers in an incentivized way. The exact setup is discussed in Section 5 where we discuss alternative behavioral forces. We ended the experiment with an ex post questionnaire. Besides background characteristics, we included Likert type statements to measure preferences for control, power, authority and reciprocity. To complement the incentivized measure of part IV we also included 10 items from Lundquist, Ellingsen, Gribbe, and Johannesson (2009) to measure attitudes towards lying.

Wave 2

Our second wave differed in four respects from wave 1. First and foremost, besides sessions with medium interest alignment (as before), we now had sessions with weak interest alignment $q = 0.4$ (instead of $q = 0.75$ in wave 1). Together the two waves thus cover a broad range of interest alignment levels.

Second, we adjusted the message space under Centralization. In half of the sessions we restricted the worker to just make factual statements about the state, i.e. to use exactly one of the following five messages: "The state is t" for $t = 1, 2, 3, 4$ and "I provide no information about the state". In the other sessions we restricted the worker's messages to recommendations only, i.e. to use one of the following four messages: "I recommend project P" for $P = A, B, C$ and "I make no recommendation". We did so to be able to verify robustness. Intuitively, workers may be more reluctant to lie if only factual statements can be made than if only (or also, cf. wave 1) recommendations can be made. Comparing the two cases may thus shed more light on the role and scope of lying aversion. Since in wave 1 subjects could use both types of messages, for $q = 0.6$ we can also investigate whether having the sheer choice between making factual statements and making recommendations matters. To keep the sessions in wave 2 otherwise comparable – and because no order effects were actually found in wave 1 – we now kept the order of games in parts I and II constant. Subjects always started with the Delegation game in part I and played the Centralization game in part II, before obtaining the choice in part III.

The final difference concerned splitting part III into two sub-parts. In the first 10 periods (part IIIa) the manager made the standard choice between centralization and delegation, with the regular option to restrict the worker's choice set under delegation. In the final 10 periods (part IIIb) the manager could only choose between centralization and *full* delegation; here the option to restrict the worker's choice set under delegation was thus absent. This change was implemented to allow for a complementary, more direct test of H3b.

Apart from the above differences, sessions in wave 2 were the same as before. The number of observations and subjects' earnings were also quite similar to wave 1. (Obviously, subjects who already participated in wave 1 could not participate in wave 2.) Overall we had another 236 subjects in wave 2. Again sessions lasted around two hours, with average earnings now equal to 25.6 euros (the minimum was 18.7 euros and the maximum 30.7 euros).

4 Results

Because the experiment was conducted in two waves, the pooled comparisons across strong, medium, and weak interest alignment partly combine variation in q with implementation differences (cf. Table 2). We therefore report the pooled three-treatment tests as evidence on the overall comparative statics, but always complement them with the corresponding bilateral within-wave comparisons: strong versus medium alignment in wave 1 and medium versus weak alignment in wave 2, with the medium-alignment treatment serving as the link across waves. These checks help separate the comparative-statics pattern from wave-specific implementation effects.

4.1 Delegation

First we consider the results under delegation when exogenously imposed. Recall that in the relevant part (either part I or II, cf. Table 2) managers and workers interact for 20 periods. To account for learning effects we focus on the decisions taken in the last 10 periods. Following the theoretical predictions in Proposition 1, we distinguish three types of outcomes based on the number of projects the manager allows the worker to choose from. Table 3 summarizes the results.

Table 3: Delegation decision manager

	$q = 0.75$	$q = 0.6$	$q = 0.4$	p-value
Del-(3)	66%	14%	13%	
Del-(2)	29%	64%	45%	< 0.01
Del-(1)	5%	23%	42%	
# of allowed projects	2.6	1.9	1.7	< 0.01

Notes: Distribution of managers' delegation decisions and average number of allowed projects conditional on the level of interest alignment q. The upper p-value comes from a permutation test (10,000 permutations) based on a 3×3 Chi-squared test statistic. The lower p-value is from a Cuzick test for trend. All statistical tests are performed on the distribution of matching-group averages.

Managers restrict the worker's choice set least frequently under strong interest alignment and most frequently under weak alignment.[14] The distribution over the three delegation outcomes differs significantly across the three treatments. This follows from a permutation test based on a 3×3 Chi-squared test statistic, using matching-group averages as the unit of observation (10,000 permutations, $p < 0.01$).[15] Another way to illustrate the exact same finding is to look at the average number of allowed projects. This equals 2.6 under strong interest alignment and reduces to 1.9 and 1.7 under medium and weak alignment, respectively. Also this trend is highly significant (Cuzick test for trend,

[14] In almost all of the Del-(2) cases managers restrict the worker's choice set to $\Delta = \{A, B\}$. In almost all of the Del-(1) cases this is either $\Delta = \{A\}$ or $\Delta = \{B\}$.

[15] For each matching group i, let $d_{1,i}$, $d_{2,i}$, and $d_{3,i}$ denote the fractions of Del-(1), Del-(2), and Del-(3) observations, respectively, with $d_{1,i} + d_{2,i} + d_{3,i} = 1$. Let $\overline{d}_j^t$ denote the average fraction of delegation outcome $j \in \{1, 2, 3\}$ among matching groups in treatment $t \in \{\text{strong}, \text{medium}, \text{weak}\}$, and let $\overline{d}_j$ denote the corresponding average across all matching groups. The test statistic used then equals $T_{\chi^2} = \sum_t n_t \sum_{j=1}^3 \frac{\left(\overline{d}_j^t - \overline{d}_j\right)^2}{\overline{d}_j}$, where n_t is the number of matching groups assigned to treatment t. The permutation test randomly reassigns treatment labels across matching groups 10,000 times and records how often the resulting statistic is at least as large as the observed value, $T_{\text{obs}} = 15.879$. This never occurs, implying $p < 1/10{,}000$; in the tables we consistently report such low values as $p < 0.01$. Analogous bilateral permutation tests, based on the corresponding 3×2 version of the same statistic, also yield $p < 0.01$ for strong versus medium alignment in wave 1, and for medium versus weak alignment in wave 2.

$p < 0.01$).[16] These comparative statics findings support Hypothesis H1.

Result 1. Managers put less restrictions on the worker's choice set when the alignment of interests becomes stronger.

Moving to the worker's project choices (unreported), in the far majority of cases he opts for the project that maximizes his payoffs given the state and his allowed choice set. In conflict state $t = 4$, the worker chooses project C under full delegation in about 75-87% of the cases, even though this project really hurts the manager in this situation. In the remaining 13-24% of the cases the worker chooses compromise project B. This happens more often under medium and weak interest alignment. A potential reason might be that full delegation is a stronger signal of trust the lower q is, leading to a stronger reciprocal reaction. If restricted to either project A or B, then in the majority of cases the worker chooses the project that maximizes his payoff given the state. By and large, under delegation managers and workers behave reasonably well in line with standard theoretical predictions.

4.2 Centralization

We next turn to (exogenously imposed) centralization, again focusing on the decisions taken in the last 10 periods out of 20 (of either part I or II, cf. Table 2). The main tests concern H2a and H2b: whether workers' messages become more informative and more influential as interest alignment increases. To interpret the size of the observed associations, we also compare them with benchmark values implied by the communication patterns characterized in Proposition 2. Because the relevant equilibrium benchmarks depend on q, these level comparisons should not be read as one maintained point prediction across all treatments.

For now we discuss workers' messages and managers' responses in isolation; in the Discussion section we verify whether workers and managers best respond to each other.

4.2.1 Workers' messages

To assess the *informativeness* of workers' messages, we measure for each matching group the association between state and message using the bias-corrected version of Cramer's V, ranging between 0 and 1 (cf. Bergsma (2013)).[17] The table below reports the averages over the different matching groups, together with the relevant statistical tests and the relevant theoretical benchmarks.

[16] Pairwise comparisons conducted separately for each wave display a similar pattern. In wave 1, the mean number of allowed projects significantly decreases ($p < 0.01$, ranksum test) when going from strong to medium alignment (2.6 versus and 1.9). In wave 2, the mean number of allowed projects under medium and weak alignment is 1.9 and 1.7, respectively; this difference is not significant ($p = 0.13$, ranksum test).

[17] We use the bias-corrected version of Cramer's V to account for the small number of independent matching group observations in our experiment. (Recall that all our statistical analyses are performed at the matching group level.) Actual and benchmark values are always computed on fixed category spaces to ensure comparability. In particular, informativeness uses the 4×4 state-message table, with states $t \in \{1, 2, 3, 4\}$ and message categories $\{M1A, M2B, M3C, M4N\}$.

Table 4: Informativeness of workers' messages

	$q = 0.75$	$q = 0.6$	$q = 0.4$	p-value
Actual Informativeness	0.74	0.70	0.67	< 0.01
Benchmark Informativeness:				
Com-(3-mixed)	**0.78**[a]	**0.70**	0.68	
Com-(2)	**0.56**[a]	**0.56**[a]	0.56[a]	
Com-(1)	0[a]	0[a]	**0**[a]	

Notes: Informativeness is based on the bias-corrected version of Cramer's V measuring the association between state and message, computed on the full 4×4 state-message category space. Benchmark rows report the value implied by the indicated communication pattern, retaining unused message categories as zero-frequency cells. The reported p-value follows from a Cuzick test for trend. Equilibrium benchmarks, characterized in Proposition 2 as a function of q, are shown in bold. Significant differences between actual and benchmark informativeness based on a Wilcoxon signed-rank test are indicated with [a] for $p < 0.01$, [b] for $p < 0.05$, and [c] for $p < 0.1$. All statistical tests are performed on the distributions of matching group averages.

Table 4 shows that the average informativeness of workers' messages significantly decreases when interests become less aligned (Cuzick test, $p < 0.01$), from 0.74 under strong alignment, to 0.70 and subsequently 0.67 under medium and weak alignment.[18] This finding is in line with Hypothesis 2a.

Result 2a. Workers' messages are significantly more informative the higher the level of interest alignment is.

Table 4 also compares actual informativeness with the benchmark levels implied by the communication patterns in Proposition 2. Under strong interest alignment, actual informativeness lies *strictly* between the Com-(2) and Com-(3-mixed) benchmarks, both of which are equilibrium benchmarks for this value of q (and thus in bold in the table). The same type of equilibria apply under medium

Influentiality (studied in the next subsection) uses the 4×3 message-action table, with the same four message categories and actions $\{A, B, C\}$. Benchmark values retain unused off-path categories as zero-frequency cells, and are therefore not conditional on the active equilibrium support.

[18] Pairwise comparisons conducted separately for each wave show similar results. In wave 1, the level of informativeness under strong and medium alignment is 0.74 and 0.69, respectively, and this difference is significant ($p = 0.02$, ranksum test). In wave 2, the level of informativeness under medium and weak alignment is 0.71 and 0.67, respectively, and this difference is significant ($p < 0.01$). Informativeness under medium alignment does not significantly differ across waves ($p = 0.27$).

alignment, where actual informativeness is very close to Com-(3-mixed) and clearly above the Com-(2) equilibrium benchmark. Under weak alignment, the relevant equilibrium prediction is Com-(1), and actual informativeness is significantly higher than this benchmark (and higher than Com-(2) as well). In fact, the observed level is then close to the descriptive Com-(3-mixed) benchmark. Thus, while the comparative statics in informativeness are in line with H2a, workers communicate more informatively than the equilibrium point prediction in the weak-alignment treatment, and the observed levels more generally resemble relatively informative communication patterns roughly similar to Com-(3-mixed).

To investigate this further we consider in Figure 1 a detailed overview of worker's messages by state and treatment. Recall that (in wave 1) workers could formulate their messages either in terms of making a claim about the state, or in terms of making a recommendation to the manager which project to choose. In the figure we bunch these two different types of messages together, as follows: M1A reflects messages that either claim that $t = 1$ or that recommend project A, and (mutatis mutandis) similarly so for M2B and M3C. M4N contains the messages that either claim that conflict state $t = 4$ applies, or communicate "I provide no information about the state", or communicate "I make no recommendation". The latter two messages arguably also concede (albeit more implicitly) that the conflict state pertains.

Figure 1: Workers' actual messages by state and treatment

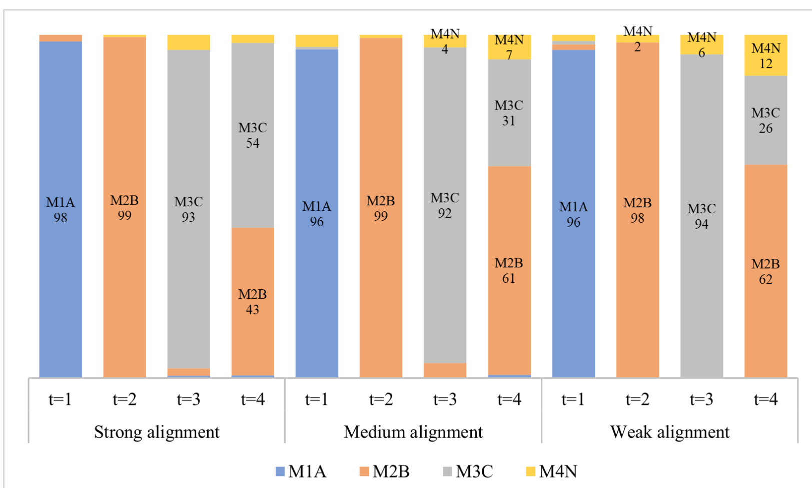

Note: M1A reflects messages that either claim $t = 1$ or recommend project A, M2B that either claim $t = 2$ or that recommend project B, M3C that either claim $t = 3$ or that recommend project C, and M4N messages that either claim $t = 4$ or make "No recommendation". Each bar gives the percentage-wise distribution over these different categories of messages (with the numbers within the bars reflecting the percentages of the largest categories).

In the three congruent states 1, 2 and 3, workers predominantly announce the true state truthfully or otherwise recommend the project that is best for both. In state 4 where interests are in conflict, however, workers tell an explicit lie in the far majority of cases. Only in 2%, 7%, and 12% of the cases under strong, medium and weak alignment respectively, the worker then can be considered truly open, by sending either “the state is 4” or “I provide no information about the state” / “I make no recommendation”.

For all treatments, actual worker behavior appears qualitatively in line with the Com-(3-mixed) benchmark. This benchmark prescribes truthful messages in the three congruent states and mixing between M2B and M3C in the conflict state. The aggregate frequency distributions over messages in state $t = 4$ do not closely track the equilibrium mixing probabilities, however. For $q = 0.75$ the latter equal $\tau_2 = \frac{1}{11}$ and $\tau_3 = \frac{10}{11}$, while for $q = 0.6$ these become $\tau_2 = \frac{6}{11}$ and $\tau_3 = \frac{5}{11}$. (Recall that for $q = 0.4$ the mixed equilibrium does not exist.) Actual behavior is such that under strong alignment workers choose M2B in about 43% of the cases where the conflict state occurs and M3C in about 54% (cf. Figure 1). Compared to the mixed equilibrium, they thus send message M3C substantially less often than they should (91%). In case of medium interest alignment they choose M3C in about 31% of the cases, again substantially below the 45% predicted. Comparative statics, however, are in line with theoretical predictions. Workers send M3C substantially more often in conflict state 4 when interests are strongly aligned than when interests are medium aligned.[19] Although under weak alignment the mixed strategy equilibrium does not exist, also there worker behavior is qualitatively in line with the Com-(3-mixed) benchmark (and captured better than by the Com-(1) equilibrium predicted). In the Discussion section we return to the intuitive nature of the worker's message strategy in Com-(3-mixed) that may explain its apparent drawing power.

4.2.2 Managers' project choices

Based on the message received from the worker, the manager chooses a project. We investigate the extent to which the worker's message influences the manager's project choice. In particular, we assess the *influentiality* of messages by measuring for each matching group the association between message received and project chosen, again using the bias-corrected Cramer's V (which again ranges between 0 and 1). The top panel in Table 5 reveals that – in line with Hypothesis H2b – managers are significantly more likely to choose different projects in response to different messages the larger the interest alignment is.[20]

[19] We also observe message switching at the individual level. Under strong (medium) interest alignment, about 35% (60%) of the workers who observe state 4 more than once choose different messages in these instances. This is consistent with individual-level mixing, although it need not imply deliberate randomization. Moreover, managers may effectively be exposed to such stochasticity more often, as they are each period randomly matched with a worker in the matching group. Around 60% (84%) of the managers are exposed to workers who switch messages under strong (medium) alignment.

[20] Pairwise comparisons separately within each wave yield similar results. In wave 1, the level of influentiality under strong and medium alignment is 0.75 and 0.70, respectively, and this difference is marginally significant ($p = 0.10$). In wave 2, the level of influentiality under medium and weak alignment is 0.61 and 0.41, respectively, and this difference is significant ($p < 0.01$). Influentiality under medium alignment is significantly higher in wave 1 than in wave 2 ($p = 0.03$).

Result 2b. Managers are significantly more likely to be influenced by workers' messages the higher the level of interest alignment is.

Table 5: Influentiality of workers' messages

	$q = 0.75$	$q = 0.6$	$q = 0.4$	p-value
Actual Influentiality	0.75	0.65	0.41	< 0.01
Benchmark Influentiality:				
Com-(3-mixed)	$\mathbf{0.82}^b$	$\mathbf{0.82}^a$	0.82^a	
Com-(2)	$\mathbf{0.69}^a$	$\mathbf{0.69}^b$	0.69^a	
Com-(1)	0^a	0^a	$\mathbf{0}^a$	

Notes: Influentiality is based on the bias-corrected version of Cramer's V measuring the association between message and project choice, computed on the full 4×3 message-action category space. Benchmark rows report the value implied by the indicated communication pattern, retaining unused message and action categories as zero-frequency cells. The reported p-value follows from a Cuzick test for trend. Equilibrium benchmarks, characterized in Proposition 2 as a function of q, are shown in bold. Significant differences between actual and benchmark influentiality based on Wilcoxon signed-rank test are indicated with ^{a} for $p < 0.01$, ^{b} for $p < 0.05$, and ^{c} for $p < 0.1$. All statistical tests are performed on the distributions of matching group averages.

The bottom panel of Table 5 reports the levels of influentiality implied by the three benchmarks Com-(3-mixed), Com-(2), and Com-(1), with equilibrium benchmarks again shown in bold. Under strong interest alignment, observed behaviour lies strictly between the predictions for Com-(3-mixed) and Com-(2); under medium and weak alignment, observed influentiality lies strictly between the values for Com-(2) and Com-(1). Theoretically, however, only the Com-(1) equilibrium exists in the weak-alignment case. Thus, as with informativeness, the overall comparative statics are consistent with the theoretical predictions (cf. Result 2b), but messages are more influential than predicted when interests are weakly aligned. It is important to note, however, that influentiality is a descriptive association measure rather than a payoff-based measure. A higher value means that managers condition their project choices more strongly on the worker's message; it does not by itself imply that the induced choices are more profitable or more often ex post optimal for the manager. We return to the payoff implications below.

Figure 2 sheds further light on this by providing a detailed overview of managers' project choices in response to the various categories of worker messages.[21] These can be roughly summarized as follows. After M1A the manager (almost) always implements project A. After M2B she is also very likely to follow by choosing project B. When interests become less aligned, however, she then increasingly opts for project A (2% under strong, 9% under medium, and 38% under weak alignment), but still (almost) never for project C. In case M3C is received, the manager's response is more mixed. Under strong alignment she is then about equally likely to follow by choosing C, or to opt for project A instead. Her response to M3C shifts to less of C and more of A (and also more of B) as interests diverge more.

Figure 2: Managers' actual project choices by message and treatment

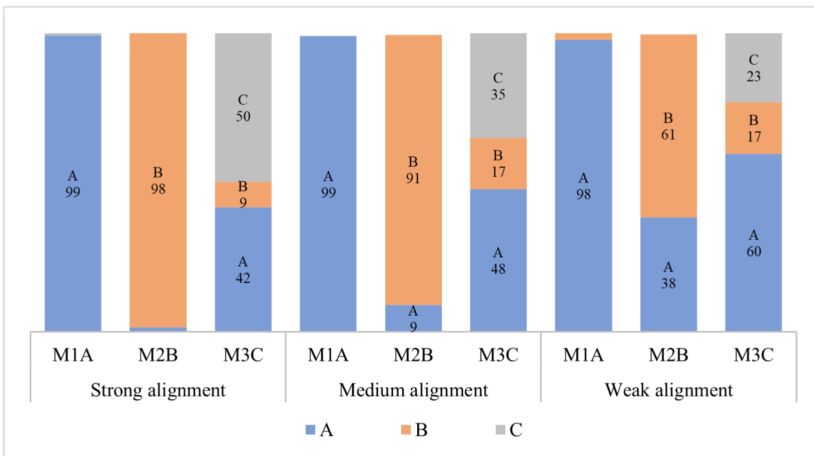

Note: Each bar gives the percentage-wise distribution over project choices A, B and C in response to the different message categories M1A, M2B and M3C (with the numbers within the bars reflecting the actual percentages for the more frequent choices).

From these general patterns the observed shifts in influentiality with interest alignment, as well as how these compare to the theoretical equilibrium benchmarks, can be intuitively understood. Influentiality equals 0 if the manager always chooses the same project, irrespective of the message received. This applies to the Com-(1) benchmark in which always project A is chosen. Influentiality increases the more often project B (in response to M2B) and the more often project C (in response to M3C) are chosen. In Com-(3-mixed) the manager chooses project C in response to M3C with a probability of 64% for any level of interest alignment q for which this equilibrium exists, so under both our strong and medium alignment treatment. The actual percentages equal 50% and 35%, respectively. This explains why influentiality falls short of this benchmark and decreases when moving from strong to medium. The manager's actual response strategy in these treatments is thus more sensitive to variations in q

[21] Managers on average receive M4N in only 2%, 4% and 9% of the cases under strong, medium, and weak interest alignment respectively. Because messages in the M4N category are rarely received, it is difficult to draw reliable inferences about the manager's response. We therefore do not include this message category in the figure.

than the mixed strategy equilibrium predicts. Under weak interest alignment, actual behavior clearly deviates from the predicted Com-(1) benchmark, as projects B and C are still chosen by the manager with non-negligible probabilities. In that respect their behavior then also shows some similarities with the Com-(3-mixed) benchmark, although managers are under weak alignment also more likely to mix if they receive M2B.[22] Whereas under weak alignment the worker's message pattern remains close to the descriptive Com-(3-mixed) benchmark, managers' responses in fact shift in the direction of the behavioral-type benchmark Com-(3A-mixed) discussed in Section 5.

4.3 Choice between centralization and delegation

Before turning to managers' endogenous organizational choices in part III, we use the exogenously imposed centralization and delegation settings from parts I and II to investigate whether the observed use of the worker's private information indeed translates into payoff-relevant gains for the manager. This is not automatic: under centralization, higher informativeness and influentiality need not coincide with better decisions, as the manager may be strategically misinformed; under delegation, greater discretion need not improve the manager's payoff if workers use it opportunistically. Appendix C therefore compares both organizational forms with the relevant no-information outcome, in which a manager would always choose project A (since $q < 4/5$ in all treatments). Expressed as improvements relative to this outcome, the predicted gains from using the worker's private information are larger when interests are better aligned. Empirically, both centralization and delegation indeed generate larger payoff improvements relative to the no-information benchmark when interests are more strongly aligned. Under weak alignment, however, these payoff gains are negative, reflecting that always choosing project A is then already quite attractive for the manager. (The complementary managers'-preferred-project measure displays the same monotone pattern but is positive only under strong alignment.) Comparing organizational forms, centralization yields significantly larger improvements only under medium alignment; under strong and weak alignment differences are not significant. These payoff comparisons provide the relevant background for the organizational-choice analysis below.

In part III of the experiment, the manager chooses whether to delegate the project choice to the worker, or to keep decision authority and select the project herself. Theoretically this choice of

[22] Around 56% of the managers under strong alignment that receive M3C more than once respond differently on these multiple occasions. For managers under medium and weak alignment this happens in 44% and 56% of the cases. This points at some individual level mixed play. At the matching group level, 75%, 90%, and 78% of the workers under strong, medium and weak interest alignment, respectively, are exposed to a mixed strategy in response to M3C, i.e. perceive mixed play by the manager. In the weak alignment treatment 57% of the managers also respond differently to M2B on the multiple occasions they observe this message, while under strong and medium alignment this only happens in 7% and 21% of the cases. At the matching group level, 7%, 32%, and 89% of the workers under strong, medium and weak interest alignment, respectively, are exposed to a mixed strategy in response to M2B, i.e. perceive mixed play by the manager.

organizational design can go either way, leading to our Hypothesis H3a that opting for delegation *per se* exhibits no systematic monotone relationship with the level of interest alignment. To investigate this we focus on the first 10 periods of the third part (i.e. part IIIa), as managers in both waves then face the same endogenous choice of C versus D (cf. Table 2).

Table 6: Choice between centralization and delegation

	C versus D (Part IIIa, both waves)			C versus full D (Part IIIb, wave 2 only)	
	$q = 0.75$	$q = 0.6$	$q = 0.4$	$q = 0.6$	$q = 0.4$
Centralization	**63%**	**65%**	**71%**	**85%**	**94%**
Delegation	**37%**	**35%**	**29%**	**15%**	**6%**
- Del-(3)	29%	3%	1%	15%	6%
- Del-(2)	7%	27%	16%	n.a.	n.a.
- Del-(1)	1%	4%	13%	n.a.	n.a.

Notes: Distribution of managers' choices between centralization and delegation by level of interest alignment. For delegation we provide a subdivision indicating the number of allowed projects in part IIIa (first 10 periods of part III), where the manager can restrict the worker's discretion under delegation. In part IIIb of wave 2 (last ten periods), the delegation choice always implies full delegation where all projects are allowed, i.e. Del-(3).

The data reveal that over these first 10 periods, managers delegate in 37% of the cases under strong alignment, 35% under medium, and in 29% of the cases under weak alignment. The Cuzick test for trend is insignificant ($p = 0.17$).[23] Thus, we do not find a significant monotone relationship between the delegation rate and the level of interest alignment. This finding is consistent with Hypothesis H3a, while leaving open the possibility of smaller treatment differences that the test may not be powered to detect. The choice of organizational design can be partially explained by managers' realized profits in parts I and II. Managers who in these previous parts received a higher profit under delegation than under centralization, are more likely to delegate in part III.[24]

[23] Results are similar if we consider the last 10 periods (i.e. Part IIIb) in wave 1. No significant differences between strong and medium alignment are then found based on a Mann-Whitney rank-sum test ($p = 0.56$). Recall that wave 1 only considered strong and medium alignment (cf. Table 2).

[24] This conclusion is based on the Spearman correlation between delegation fractions and the relative profits of delegation versus centralization in the first two parts. Under strong alignment this correlation equals 0.60, while in the medium and weak treatments it equals 0.45 and 0.34 respectively. In all cases it is significant at the 1% (or better) level. No significant correlations are found with items in the ex post questionnaire that relate to control, power and

Result 3a. The manager's choice of delegation per se exhibits no significant monotone relationship with the level of interest alignment.

Although delegation fractions exhibit no significant monotone trend when interests become better aligned, the nature of delegation (if chosen) does change. That is, similar to the case where delegation is exogenously imposed (cf. Table 3), there is a shift at the intensive margin by giving the worker more freedom of choice. In particular, the average number of allowed projects now equals 2.7 under strong, 2.0 under medium, and 1.6 under weak interest alignment.[25] Also here the Cuzick test for trend is highly significant ($p < 0.01$).[26]

Together, the absence of a significant monotone relationship in delegation *per se* (cf. Result 3a) and the observation that the worker's discretion under (endogenous) delegation significantly increases with q, suggest that full delegation may become more likely when interests become better aligned. Part IIIb of wave 2 provides a complementary, more direct test of this. In these final 10 periods of part III, the manager could only choose between centralization and *full* delegation (cf. Table 2). Under the latter the worker keeps the full set of options of implementing either project A, B, or C. A clear limitation of this more direct test is that only the weak and medium alignment treatments can be compared, since the strong alignment treatment was not considered in wave 2 (see again Table 2). The observed frequency of full delegation equal 6% under weak alignment and 15% under medium alignment. Although this increase is in line with theoretical predictions (cf. Hypothesis H3b), it is not significant according to a Mann-Whitney rank-sum test ($p = 0.20$).[27] When looking at the part IIIa data of wave 2, the full delegation fraction (i.e. outcome Del-(3)) equals 5% under medium alignment and 1% under weak alignment. This difference is significant according to a Mann-Whitney rank-sum test ($p = 0.03$).[28] Including the strong alignment treatment in this comparison as well, a Cuzick test

authority.

[25] Relative to Table 3, the distributions over Del-(3), Del-(2) and Del-(1) are somewhat more pronounced: 75%, 24% and 1% for $q = \frac{3}{4}$, 11%, 77% and 12% for $q = \frac{3}{5}$, and 1%, 56% and 43% for $q = \frac{2}{5}$. Worker behavior is similar as under exogenous delegation; he again typically chooses the project that maximizes his payoffs given his allowed choice set and the state.

[26] Part III behavior under endogenously chosen centralization is also closely in line with observed behavior when it is exogenously imposed. The worker again almost always sends message Mt in the congruent states $t \in \{1A, 2B, 3C\}$. In conflict state $t = 4$, he mixes between M2B and M3C; under strong alignment with about equal probabilities, while under medium and weak alignment M2B is chosen substantially more often than M3C is (like before). The worker is thus equally willing to share information as under exogenous communication, even though now delegation was an explicit option but forfeited by the manager. Also the observed pattern of managers' behavior is very similar. She almost always follows M1A by implementing project A. If M2B is received, then in the far majority of cases she implements project B under strong and medium alignment, while under weak alignment she mixes between A and B. Finally, after M3C the manager mixes mostly between projects A and C; the latter is chosen in 54%, 42%, and 33% of the cases under respectively strong, medium, and weak alignment.

[27] The delegation fractions in part IIIa of wave 2, where the manager can restrict the worker's choice set, equal 29% under weak and 35% under medium alignment. These fractions differ significantly from the fractions of full delegation in part IIIb for the corresponding treatment (Wilcoxon sign-rank test, $p < 0.01$ in both cases). The manager is thus significantly less likely to delegate per se when she cannot restrict the worker's discretion.

[28] In wave 2, both under medium and under weak alignment the full delegation (Del-(3)) fraction in part IIIa is lower than in part IIIb; the difference is significant only for the weak alignment treatment (p-value under medium (weak) alignment based on Wilcoxon sign-rank test equals $p = 0.16$ ($p < 0.01$)). Similar results are obtained when for medium

reveals a significant downward trend in the full delegation fraction as interests become less aligned ($p < 0.01$), based on the combined data from both waves. Because we do not have data for $q = 0.75$ in part IIIb, we cannot test for such a trend in that part. The overall statistical evidence for significant differences is therefore quite fragile.

Result 3b. The manager's choice between *full* delegation and centralization is not significantly different under medium and weak alignment.

Overall, the evidence in favor of Hypothesis 3b is thus more limited. Result 1 and Result 3a together provide supportive evidence, while Result 3b does not.

5 Discussion

The observed comparative statics results reported in the previous section are largely in line with our hypotheses derived from standard equilibrium predictions under expected payoff maximization. At the same time, subjects' behavior also clearly deviates from the equilibrium point predictions, especially so under centralization in the weak interest alignment treatment. In this section we consider various behavioral forces that may shift the actual tradeoff between (optimal) delegation and centralization.

5.1 Delegation

A prominent behavioral force that in principle could affect the outcome under delegation is worker reciprocity. Negative reciprocity may motivate the worker to purposely make a suboptimal project choice if the manager intentionally restricts his choice set. Such hidden costs of control have been observed in related principal-agent settings with no private information (cf. Falk and Kosfeld (2006)). In our setting this could potentially occur only if the worker is restricted to $\Delta = \{A, B\}$. The most salient instance for the worker to punish then occurs when $t = 3$; in this state the manager's restriction to $\Delta = \{A, B\}$ really hurts both.[29] Actual worker behavior is such that, if restricted to either project A or B and state 3 materializes, the worker chooses A in only 3% (under weak) up to 10% (under strong) of the cases. We thus do not find substantive hidden costs of control in our setting. In a similar vein, positive reciprocity could potentially motivate the worker to reward the manager for not restricting his discretion. Full delegation signals trust and can be considered a kind act, to which a reciprocal worker might respond by using his authority more "responsibly". In particular, the worker may implement project B as a compromise in conflict state $t = 4$ if authority is fully delegated. (And

alignment both waves are included in the analysis.

[29] Note that in conflict state $t = 4$, a worker deviating from his own best project B (towards implementing A) would in fact be rewarding the manager.

in case the worker would indeed always do so, the principal would not want to limit his discretion as long as $q > 3/8$.) The actual data show that also positive reciprocity plays a limited role; under strong, medium and weak alignment, the worker chooses the compromise project B in 13%, 24%, and 22% of the $t = 4$ cases, respectively (cf. Section 4.1).

Overall, neither the hidden costs of control nor the hidden benefits of autonomy appear to significantly influence (optimal) delegation, and thus the tradeoff with centralization.

5.2 Centralization

Under centralization, two well-established behavioral forces may help explain deviations from the benchmark equilibrium predictions: workers' lying aversion and managers' credulity. The aim of this subsection is not to identify a structural behavioral model of these forces, but to assess whether limited departures from the fully strategic benchmark can plausibly rationalize the remaining communication patterns. This distinction is important because truthful messages in the congruent states, and managers sometimes following messages at face value, may also arise under fully strategic equilibrium behavior (cf. Proposition 2). We therefore first examine whether workers and managers best respond to each other in the centralization game. We then discuss the extent to which lying aversion and credulity may account for the residual deviations.

5.2.1 Best response analysis

First we determine workers' best responses. We take managers' average behavior in a matching group as the manager's actual strategy (like before we focus on the final 10 periods of the relevant parts). For each state we subsequently determine which message would maximize the worker's payoffs against this actual strategy and we label this as his 'best response message'. Table D.1 in Appendix D provides a detailed overview of these best responses given the state and the level of interest alignment. The overall pattern yields two main observations. First, in all three congruent states the worker is (almost) always fully honest and this (almost) always coincides with best responding. Second, in conflict state $t = 4$ workers typically choose either M2B or M3C (and either of these quite often), and also best responses always correspond to one of these two messages.[30] Yet workers' actual choice frequently deviates from their best response; that is, M2B is often chosen in instances where M3C would be best and vice versa. As shown in Figure 3a, under strong and medium interest alignment workers overall send M3C somewhat more often and M2B somewhat less often than would be optimal given

[30] As noted before in Subsection 4.2, the worker sends message M4N in only 2% (strong), 7% (medium) and 12% (weak) of the cases. These numbers are too low to make reliable inferences at the matching group level about being a best response or not. The aggregate distribution of manager's project choices over all matching groups, however, clearly suggests that they are not and M3C would have been better for all levels of interest alignment considered (cf. Table D.2 in Appendix D).

managers' behavior. In case of weak alignment this is the other way around.

Figure 3: Workers' and managers' best responses by treatment

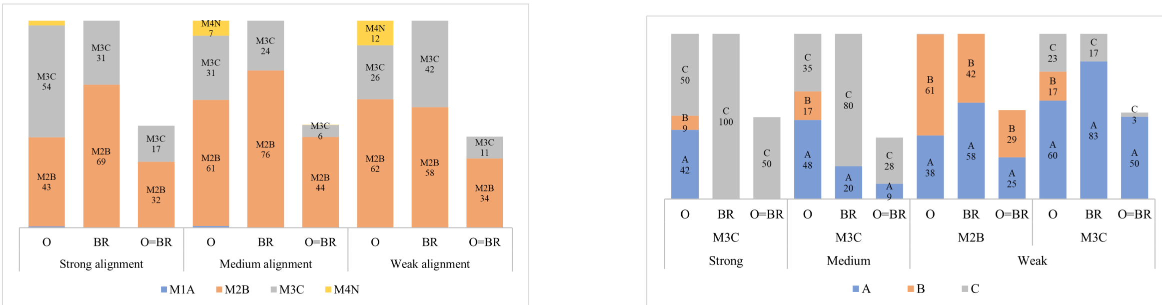

(a) Workers' best response messages if state is 4

(b) Manager's best responses by messages

Note: Panel (a) presents the worker's behavior if the worker observes state 4 and panel (b) presents the manager's behavior if M3C is received (for $q = 0.4$ we also present managers' choices if M2B is received). Bars labelled "O" reflect observed behavior, bars labelled "BR" reflect best responses, and bars labelled "O=BR" reflect observed behavior that coincides with best response.

Second, in a similar vein as for workers, we determine the extent to which managers are best responding to workers' average behavior in a matching group. We take workers' average behavior in a matching group as the worker's actual strategy (like before we focus on the final 10 periods of the relevant parts). For each message we subsequently determine which project would maximize the manager's payoff against this actual strategy and we label this as her 'best response project'. Table D.2 in Appendix D reports for each treatment managers' best responses conditional on the message received. This overview reveals that managers always best respond to message M1A by choosing project A and, under both strong and medium alignment, also almost always best respond to M2B by choosing project B. Deviations from best responses are larger when M2B is received in the weak alignment treatment. In that case managers often opt for project B, while project A is more often the best response. Nevertheless, managers' alternating behavior between frequently choosing A but also often choosing B is – although clearly not in the right proportions – qualitatively in line with best response behavior; also the latter alternates between often coinciding with A and often with B (but never with project C). In case M3C is received, managers' deviations from best response are larger in all treatments. Under both strong and medium alignment, managers then frequently opt for project A while project C is either always (strong) or often (medium) the best response. And under weak alignment, the manager opts for project B in 17% of the cases, while the best response is to either choose A (most often) or C (in a minority of instances). Figure 3b visualizes the discussed deviations. A main overall takeaway from the managers' best response analysis is that when interest alignment decreases, there is a clear shift to project A both being chosen and being a best response more often, not only after M3C but also after M2B in case interest alignment is weak.

A potential reason why workers in the experiment always best respond by being fully honest in the congruent states is that there is a modest fraction – denoted by $\gamma \in [0, 1]$ – of credulous managers that always believe and follow the worker's message. To intuitively see this, note that equilibrium Com-(1) in Proposition 2 is not robust to having $\gamma > 0$. In this equilibrium workers always send message m_1 and obtain project A in response. Workers do so because other messages would necessarily lead to project A being implemented as well. With just a few credulous managers, however, the latter no longer holds. A strategic worker knowing that $t = 2$ can then strictly gain by deviating to sending m_2, because this message will be taken at face value by the credulous managers and thus induce project B if matched to one of them (yielding the worker an expected payoff of $100\gamma > 0$). A similar observation holds for the Com-(2) equilibrium whenever $\gamma > \frac{1}{10}$. In state $t = 3$ a strategic worker can then strictly gain by choosing m_3 instead of m_2 to obtain $100\gamma > 10$ in expected payoffs. Best responses and equilibrium predictions thus change considerably in the presence of just a few behavioral types. Theoretically, a small fraction of credulous managers shifts the distribution of equilibrium outcomes per se, just by now having these types (who behave differently) in the population; an 'extensive margin' effect. Yet their presence may potentially also have a large(r) 'intensive margin' impact, by changing the equilibrium behavior of the majority group of strategic agents.

In Appendix B we derive the equilibrium predictions in the presence of some modest fraction of credulous managers ($\gamma > \frac{1}{10}$), as well as some fraction $\lambda \in (0, 1)$ of lying averse workers. Relative to the standard equilibrium predictions presented in Proposition 2 there are two major changes. First, per the intuition above, workers necessarily always tell the truth in all congruent states, so Com-(1) and Com-(2) no longer exist. Second, the equilibrium predictions shift to project A becoming a best response for strategic managers more often. In particular, for lower levels of interest alignment now an alternative mixed strategy equilibrium – labeled Com-(3A-mixed) – exists, in which the strategic worker continues to communicate the true state in all congruent states and to mix in the conflict state. But now a strategic manager chooses project A for sure after m_3 (rather than mixing between A and C as in Com-(3-mixed)), and mixes between projects A and B in response to m_2 (rather than project B for sure). For this alternative equilibrium to theoretically apply under weak alignment instead of Com-(3-mixed), it is required that the fraction of lying averse workers is sufficiently small: $\lambda < \frac{8}{33} \approx 0.24$ should hold. Here the intuition is as follows. The strategic worker's mixing strategy in the conflict state $t = 4$ between m_2 and m_3 is fully determined by making the strategic manager indifferent between A and B in response to m_2. For this given mixing strategy, the strategic manager should also strictly prefer project A in response to m_3. The latter can only happen when the fraction of lying averse workers λ is not too high. Note that our finding that especially managers' behavior changes when going from intermediate to weak interest alignment is qualitatively in line with a move from Com-(3-mixed) to Com-(3A-mixed).[31]

In summary, a low fraction of lying averse workers and especially a modest fraction of credulous managers can have substantial intensive margin effects on the (best response) behavior of the majority group of strategic types. Theoretically this shifts behavior from the standard predictions in Proposition 2 in the direction of what we observe. We therefore next explore whether our experimental data provides some evidence for (limited) lying aversion and credulity.

5.2.2 Lying aversion and credulity

Both observed behavior under centralization in parts I to III and the incentivized measures employed in part IV may provide some evidence of having lying averse workers and credulous managers among our subjects. As discussed in the previous subsection, however, it is difficult to cleanly identify such behavioral types based on the former, given that strategic behavior often coincides with being honest or being credulous. The separate measures in part IV in principle allow for cleaner identification, but have the potential drawback that they are derived from a different strategic setting. We therefore look at both sources of information and also investigate the extent to which these correlate.

Lying Aversion First consider actual worker behavior under centralization. A sufficiently lying averse worker would want to shy away from explicitly telling a lie in the conflict state. In only 2% (strong), 7% (medium), and 12% (weak) of the $t = 4$ cases workers can be considered truly honest, as they then send message M4N. This suggests that the actual fraction of lying averse workers is very low. However, lying aversion can also result in workers preferring making a recommendation over making a factual statement in the conflict state. If messages have a literal meaning, then only for factual statements it is (arguably) clear to everyone what would be considered a lie. One would thus expect workers to be less explicit in the conflict state, essentially reverting to more vague communication as in Serra-Garcia et al. (2011).[32] The data from wave 1, in which workers could either make a factual statement or send a recommendation, can shed some light on this. In the three congruent states workers send a message about the state in 81% (80%) of the cases under strong (medium) alignment, while in the conflict state this happens in 76% (73%).[33] This decrease is marginally significant only

[31] As the discussion in Appendix B reveals, for a very large fraction of credulous managers $\gamma > \frac{7}{11}$ the mixed strategy equilibria do not exist. The intuition is that when the fraction of credulous managers is high, then workers in state 4 prefer to always send m_3. Then an equilibrium Com-(3A) exists that is equivalent to Com-(3) in terms of worker behavior. The single difference with Com-(3) is that the strategic manager then (again) shifts to choosing A (rather than C) in response to m_3. Note that this equilibrium is not in line with worker's actual behavior when $t = 4$ (as they then frequently send m_2).

[32] In a similar vein, Özer et al. (2018) argue that, compared to information sharing, advice provision results in more deliberative thinking and more attention for profit implications, resulting in more self-serving communication behavior.

[33] Table E.1 in Appendix E provides detailed information on the type of message (either factual statement or recommendation) sent by the worker in wave 1 by state and treatment.

under medium interest alignment ($p = 0.10$ under medium and $p = 0.53$ under strong according to a Wilcoxon signed-rank test). The fact that the relative fractions remain largely the same across different states suggests that subjects do not adapt the type of language they use in the conflict state considerably.

Figure 4: Workers' messages in the conflict state under medium interest alignment

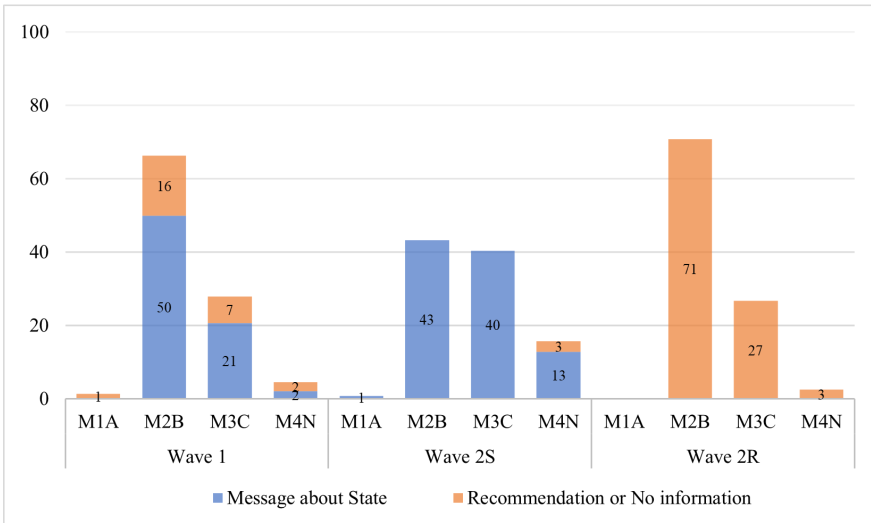

Note: Wave 1 presents the worker's messages in State 4 in wave 1. Wave 2S and Wave 2R present the worker's messages in State 4 in wave 2, where in Wave 2S (Wave 2R) the worker could only send information about state (make recommendation).

In wave 2 workers could either only send messages about the state (sessions with 'S' in Table 2), or only make a recommendation (sessions with 'R'). For medium interest alignment ($q = 0.6$) we can thus compare whether worker's behavior differs if messages are restricted to either factual statements or to making recommendations only (as in wave 2), or if both types of messages are allowed for (as in wave 1). Figure 4 depicts this comparison. In wave 1 the worker sends a truthful message in 4.5% of the cases in which the conflict state applies. When the worker can only communicate about the state a truthful message is sent in 15.7% of the conflict state cases, while with recommendations only this reduces to 2.5%. The worker thus sends a truthful message significantly more often when only factual messages are possible than when making a recommendation is also an option ($p = 0.03$, Mann-Whitney ranksum test), or when making a recommendation is the only possibility (again $p = 0.04$). Also under weak alignment ($q = 0.4$) the worker sends a truthful message more often when only factual messages are possible (16.7%) than when only recommendations can be made (7.1%), but here the difference is not significant ($p = 0.17$).[34] Overall we thus obtain some limited evidence that workers are somewhat more reluctant to lie with factual statements. This is in line with having some (small) subset of lying

[34]Table E.2 in Appendix E shows the worker's messages in wave 2 by state and treatment.

averse workers in our sample. Alternatively, it could also derive from strategic workers expecting that especially an outright revelation of the actual conflict state will induce a more favorable response than a questionable recommendation would do. The aggregate distribution of managers' project choices, however, suggests otherwise, as strategic workers would still be better off sending M3C instead.

We next consider the incentivized measures included in part IV based on an adjusted version of Gneezy et al. (2013). Workers in the earlier parts were assigned role A and randomly matched to a former manager having role B.[35] A's obtained private information about an integer between 1 and 6 that was randomly assigned by the computer to each pair and had to send a – not necessarily truthful – message m to B about this outcome. A's earned more the higher their message, viz. $500 + 100m$. B's could either follow ('believe') A's message, then earning 500 only if the message equals the actual number assigned, or not follow and getting 150 points for sure. The strategy method was employed, asking A for her message for any possible assigned number (between 1 and 6) and asking B for his decision for any message that he might receive. Based on these strategies, the choices that resulted from the random number actually assigned to a pair were carried out.[36] Note that this is a high-conflict sender-receiver game; the single (standard) equilibrium outcome is $m = 6$ and 'not follow'. Of the overall 233 workers, 34% always reported truthfully, while 39% always reported 6. The reports of the remaining 27% were more mixed. These numbers suggest that the fraction of lying averse subjects lies somewhere in between 0.25 and 0.5, which accords reasonably well with the findings of earlier (either constant or high conflict) experiments. Yet they do not seem to match with our observation above that in the centralization game, true honesty occurs in overall less than 7% of the conflict cases. Apparently those who are averse to lying in the high conflict game of part IV have less of a problem with it in the centralization game where interests are often, but not always aligned.

A similar conclusion follows from correlating individual lying behavior in the two different parts. Based on part IV we measure 'lying proneness' by the aggregate (over all possible outcomes) extent to which a worker overstates the actually assigned number. This measure ranges from 0 (when always reporting truthfully) to 15 (when always claiming $m = 6$). For the individual measure of lying behavior in the centralization game, we calculate the proportion of conflict state cases in which the worker lies. Lying in state 4 corresponds to either sending a factual statement that the state is 1, 2, or 3, or making a recommendation for project B or C. The Spearman rank correlation between the measure of lying behavior and lying proneness appears to be small (0.10) and not significant (p-value of 0.14).[37] Overall we therefore conclude that our strategic situation makes that even inherently lying

[35] We did not explicitly inform subjects of the Worker-A and Manager-B mapping; after having read the instructions for part IV (see Appendix E), subjects just learned their respective roles in that part.

[36] The instructions did not make explicit that each of the six numbers had an equal chance of being assigned, yet the wording in the instructions – "The computer will randomly assign one of the numbers 1, 2, 3, 4, 5, or 6 to each pair" – arguably sufficiently conveys that to the subjects.

[37] An alternative measure of lying proneness follows from adding the scores on the 10 items taken from Lundquist et

averse subjects have little problem telling a lie only now and then.

Credulity Inferring managers' credulity based on their project choices in the centralization game is difficult. Although managers often do not best respond to messages M2B and M3C, they also do not deviate in a systematic way (cf. Subsection 5.2.1). For instance, although under weak alignment managers are on average somewhat too credulous when receiving message M3C, under strong and medium alignment they are actually too skeptical in their response to this message. Yet, similar to lying aversion, credulity may potentially also become apparent from different responses to factual statements and corresponding recommendations. Above we observed that workers tend to make factual statements somewhat less often in the conflict state than in the congruent states, although this pattern is not fully uniform within each message category.[38]This suggests looking directly at whether managers respond differently to the two message formats. We therefore calculate the ratio of 'following the worker's factual statement' and 'following the worker's recommendation'. We do so for each of the three relevant message categories M1A, M2B and M3C, focusing on wave 1 data because only there managers are exposed to both types of messages. The actual ratios obtained under strong (medium) alignment equal 0.99 (1.02) for M1A, 1.13 (1.28) for M2B, and 1.71 (1.54) for M3C. According to a Cuzick test this increasing trend is significant only for the strong interest alignment treatment ($p < 0.01$; $p = 0.61$ under medium). We thus find at most suggestive and non-uniform evidence that managers treat factual statements as more credible than corresponding recommendations. Managers therefore do not appear very credulous in this regard.

The lying game in part IV provides a cleaner measure of credulity. We aggregate the extent to which an individual manager follows presumably overstated reports. Intuitively, following after a message that the assigned number equals 1 is not credulous. Apart from the instance in which this is truly the case, there are no other actual outcomes for which subject A has an incentive to lie and report 1. This is different for message $m = 6$. Besides the instance in which the assigned number indeed is 6 (truth-telling), there are also five other outcomes for which A has an incentive to lie and report 6. Following after $m = 6$ is thus particularly credulous. Our overall measure of 'credulity proneness' weighs all messages accordingly. It ranges from 0 (only follow after $m = 1$) to 15 (always follow after any message). Among the 233 managers, 47% always decided to follow while 6% is not credulous at all and only believes message 1. In line with other experiments (e.g. Gneezy (2005)), we thus find that the mass of credulous prone receivers (i.e. managers) is somewhat larger than the mass of lying-averse senders (workers).

al. (2009) included in the ex post questionnaire (participants answered these items on a 1 (fully disagree) to 7 (fully agree) scale). Also for this measure the Spearman correlation is small (0.12) but now marginally significant ($p = 0.07$). The correlation between the two different lying proneness measures equals 0.13 ($p = 0.06$).

[38]When considering the factual-statement share for M1A, M2B, and M3C by treatment using the wave 1 data, the share is lower in the conflict state in five cases and slightly higher in only one case, namely M3C under strong alignment.

To investigate the extent to which 'credulity proneness' as measured in part IV can explain managers' actual behavior under centralization, we determine the proportion of cases in which the manager chooses to follow the worker's message in the centralization game.[39] As noted in the previous subsection, managers always best respond to message M1A by choosing project A. Credulity is therefore best assessed when the manager receives either M2B or M3C. We calculate the proportion of cases in which the manager chooses project B after receiving M2B or project C in response to M3C. The Spearman correlation between this proportion and 'credulity proneness' from part IV equals 0.11 and is marginally significant ($p = 0.09$).[40] A similar conclusion follows if we look at factual statements and recommendations in isolation, using wave 2 data. For factual statements either claiming that 'the state is 2' or that 'the state is 3', the correlation is 0.05 and insignificant ($p = 0.47$). For messages recommending either B or C it equals 0.18 ($p = 0.03$). Taken together, also credulity does not appear to be a strong driver of behavior under centralization.

In summary, both lying proneness and credulity proneness among our subjects appear similar to those observed in other experimental studies of sender-receiver games (Gneezy (2005); Gneezy et al. (2013)). However, these direct type measures provide at most modest support for the behavioral mechanism, as they translate only weakly into lying-averse or credulous behavior in the strategic setting of our centralization game. One explanation here might be that our design is based on messages of a fixed format. Typically, such fixed-format messages lead to more lying than free-format chat does (cf. Charness and Dufwenberg (2010)). Another factor could be that sender and receiver often have common interests in our setting: in all congruent states, the sender is best off telling the truth. This may make lying in the conflict state less problematic, as this state occurs only occasionally. This differs from experimental tests of the standard Crawford-Sobel framework, such as Cai and Wang (2006), where the alignment of interests is state independent and the sender effectively always has an incentive to portray the state differently from how it actually is. Thus, the main evidence for limited behavioral heterogeneity is not the direct type measures themselves, but rather the qualitative fit between the aggregate communication patterns and the behavioral-type benchmarks.

[39] The manager follows the worker's message if she considers statements about the state as true, follows the worker's recommendation, and decides on the basis of the prior belief if the worker makes no recommendation.

[40] A potential drawback here is that in many instances implementing project B after M2B or project C after M3C is the best response. To account for this we also consider an alternative measure where we eliminate the cases where the manager is best responding. That is, we calculate the proportion of cases in which the manager chooses project B (C) in response to M2B (M3C) while in fact another project would be the best response. The Spearman correlation between this alternative proportion and credulity proneness is small (0.10) and insignificant ($p = 0.16$).

5.2.3 Limited rationality and non-equilibrium play

Although at the aggregate level actual outcomes are qualitatively in line with the mixed strategy equilibria and subjects do not deviate from best responding in a systematic way,[41] it is still quite likely that the pattern that we observe results from out-of-equilibrium play. Arguably our centralization game is somewhat complex and subjects might differ in their level of strategic sophistication. We thus briefly explore a level-k analysis of our setting. Level-k models allow for heterogeneity between subjects and assume that level-k players L_k best respond against level L_{k-1}. Beliefs about others' choices or cognitive levels need not be correct though, making it a non-equilibrium approach. Key and degree of freedom in any level-k model is the assumption about how the non-strategic level-0 players behave (besides the distribution of the players over the different cognitive levels). In line with the above discussion of lying aversion and credulity, we assume that level-0 workers are lying averse (i.e. 'naive' in the sense of always telling the truth) and that level-0 managers are credulous (i.e. naive in always following the worker's message). The following choices for the different levels then emerge:

Table 7: Simplified level-k analysis of the centralization game

Worker	$t=1$	$t=2$	$t=3$	$t=4$	Manager	m_1	m_2	m_3	m_4
L_0	m_1	m_2	m_3	m_4	L_0	A	B	C	A
L_1	m_1	m_2	m_3	m_3	L_1	A	B	C	A
L_2	m_1	m_2	m_3	m_3	L_2	A	B	A or C	A
L_3	m_1	m_2	m_2 or m_3	m_2 or m_3	L_3	A	B	A or C	A
L_4 ff.	m_1	m_2	m_2 or m_3	m_2 or m_3	L_4 ff.	A	A or B	A or C	A

Note: Entries with "or" are q-dependent best responses. For the experimental values $q \in \{0.75, 0.6, 0.4\}$, all of which are below 33/43, the A or C entry after m_3 should be read as A. The L_3 worker row then predicts m_2, rather than m_3, in both states $t=3$ and $t=4$. The remaining A or B entry after m_2 should be read as B for $q=0.75$ and $q=0.6$, and as A for $q=0.4$, since the cutoff is 12/23.

Obviously, the above level-k analysis is highly simplified and for illustrative purposes only. In applying level-k reasoning to a related communication-delegation experiment, Lai and Lim (2012) take a much more elaborate approach by considering alternative specifications of non-strategic L_0 behavior and disciplining the implied distribution of subjects over sophistication levels. Our purpose here is more modest. We take honest workers and credulous managers as the natural L_0 benchmark and use Table 7 only to illustrate how such limited departures from fully strategic behavior may be

[41] Note that a mixed strategy equilibrium outcome does not necessarily require deliberate randomization by individual players, but may also be seen as a stochastic steady state of a large population or result from players choosing pure strategies in a slightly perturbed game (Harsanyi's purification theorem); see e.g. Osborne and Rubinstein (1994) (Chapter 3.2) for an insightful discussion of the various interpretations of mixed strategy equilibria.

amplified through higher-order best responses.

The level-k analysis in Table 7 can then be read in a similar spirit as the behavioral-type equilibrium analysis in Appendix B, although the two channels are distinct. In Appendix B, strategic players best respond to an actual population mixture of behavioral and strategic types; in the level-k exercise, higher-level players instead best respond to lower-level reasoning. Starting from honest workers and credulous managers as L_0's, higher-level workers continue to communicate truthfully in the congruent states up to a point, but behave more strategically in the conflict state. Higher-level managers, anticipating such concealment, respond by choosing project A more often after messages that are effectively ambiguous, i.e. the project that is best for them in the state the worker would like to hide. Taken literally, however, the level-k pattern does not fully rationalize the data: for the experimental values of q, the L_3 worker row predicts m_2 rather than m_3 also in state $t = 3$, whereas workers in the data overwhelmingly continue to send m_3 in that state. The exercise should therefore not be read as a fitted account of behavior. Rather, it suggests an analogous form of strategic amplification: limited honesty or credulity, whether represented as behavioral types or as lower-level reasoning, may affect the behavior of more strategic subjects and thereby help sustain informative communication in the congruent states. Unlike the pattern emphasized by Lai and Lim (2012), we do not observe systematic over-communication relative to the full set of equilibrium benchmarks considered here. Communication is clearly more informative than some of the pure-strategy benchmarks predict, especially under weak interest alignment, but this conclusion is qualified once the relevant mixed-strategy benchmarks are taken into account. Rather, the Com-(3-mixed) outcome provides a parsimonious description of much of the observed aggregate behavior under strong and medium interest alignment, while Com-(3A-mixed) qualitatively captures the main patterns under weak alignment.

6 Conclusion

Using a laboratory experiment, this paper investigates how managers choose between communication and delegation when informed workers may have different objectives. We study a simple principal–agent environment in which only the worker observes the payoff-relevant circumstances and in which the alignment of interests between manager and worker depends on this private information. In three equally likely states, preferences are aligned; in a fourth state, preferences are opposed. By varying the probability of the conflict state, the experiment tests how interest alignment affects communication, restricted delegation, and the manager's organizational choice.

The experimental results first confirm the predicted comparative statics within delegation. Managers adjust the scope of the worker's discretion to the degree of interest alignment: when interests are less aligned, they impose tighter restrictions; when interests are better aligned, they allow more freedom. At the extensive margin, however, the manager's choice between communication and delegation per se varies little with interest alignment; most of the adjustment under delegation occurs through the extent of restrictions. Under communication, messages remain informative, but selectively so. Workers typically communicate truthfully when interests are aligned, while concealing or distorting information when interests conflict. Consequently, communication conveys more information than the relevant pure-strategy equilibrium benchmarks predict, especially when interests are weakly aligned. This excess informativeness does not reflect unconditional honesty, but rather the preservation of a meaningful language in situations where interests are aligned.

The communication patterns are qualitatively consistent with a limited role for behavioral heterogeneity. In particular, a small share of lying-averse workers and a modest share of credulous managers offer one possible mechanism for strategic amplification: the presence of some managers who take messages at face value may make it worthwhile for strategic workers to preserve an informative language when interests are aligned. Thus, the results do not suggest that communication works because workers are generally honest or managers generally naive. Rather, they point to a more limited mechanism through which small departures from the standard benchmark can affect strategic communication.

References

Agastya, M., Bag, P., & Chakraborty, I. (2014). Communication and authority with a partially informed expert. *Rand Journal of Economics*, *45*, 176-197.

Aghion, P., & Tirole, J. (1997). Formal and real authority in organizations. *Journal of Political Economy*, *105*(1), 1-29.

Alonso, R., & Matouschek, N. (2008). Optimal delegation. *Review of Economic Studies*, *75*, 259-293.

Bartling, B., Fehr, E., & Herz, H. (2008). The intrinsic value of decision rights. *Econometrica*, *82*, 2005-2039.

Bartling, B., & Fischbacher, U. (2012). Shifting the blame: On delegation and responsibility. *Review of Economic Studies*, *79*(1), 67-87.

Battaglini, M., Lai, E., Lim, W., & Wang, J. T.-Y. (2019). The informational theory of legislative committees: An experimental analysis. *American Political Science Review*, *113*(1), 55–76.

Bergsma, W. (2013). A bias-correction for Cramér's V and Tschuprow's T. *Journal of the Korean Statistical Society*, *42*, 323-328.

Blume, A., DeJong, D. V., Kim, Y.-G., & Sprinkle, G. B. (1998). Experimental evidence on the evolution of meaning of messages in sender-receiver games. *The American Economic Review*, *88*(5), 1323-1340.

Blume, A., DeJong, D. V., Kim, Y.-G., & Sprinkle, G. B. (2001). Evolution of communication with partial common interest. *Games and Economic Behavior*, *37*(1), 79-120.

Blume, A., Lai, E. K., & Lim, W. (2020). Chapter 13: Strategic information transmission: a survey of experiments and theoretical foundations. In C. M. Capra, R. T. Croson, M. L. Rigdon, & T. S. Rosenblat (Eds.), *Handbook of experimental game theory* (p. 311-347). Edward Elgar Publishing.

Brandts, J., & Cooper, D. J. (2025, 03). Managerial leadership, truth-telling, and efficient coordination. *The Economic Journal*, *135*(670), 1942-1979.

Cai, H., & Wang, J. T.-Y. (2006). Overcommunication in strategic information transmission games. *Games and Economic Behavior*, *56*, 7-36.

Charness, G., & Dufwenberg, M. (2010). Bare promises: An experiment. *Economics Letters*, *107*, 281-283.

Coffman, L. C. (2011). Intermediation reduces punishment (and reward). *American Economic Journal: Microeconomics*, *3*(4), 77–106.

Crawford, V. P., & Sobel, J. (1982). Strategic information transmission. *Econometrica*, *50*, 1431-1451.

de Haan, T., Offerman, T., & Sloof, R. (2015). Money talks? an experimental investigation of cheap talk and burned money. *International Economic Review*, *56*(4), 1385-1426.

Dessein, W. (2002). Authority and communication in organizations. *Review of Economic Studies*, *69*, 611-838.

Dominguez-Martinez, S., Sloof, R., & von Siemens, F. A. (2014). Monitored by your friends, not your foes: Strategic ignorance and the delegation of real authority. *Games and Economic Behavior*, *85*, 289-305.

Evdokimov, P., & Garfagnini, U. (2019). Communication and behavior in organizations: An experiment. *Quantitative Economics*, *10*(2), 775-801.

Falk, A., & Kosfeld, M. (2006). The hidden costs of control. *American Economic Review*, *96*(5), 1611-1630.

Farrell, J. (1987). Meaning and credibility in cheap-talk games. *Games and Economic Behavior*, *5*, 514-531.

Fehr, E., Herz, H., & Wilkening, T. (2013). The lure of authority: Motivation and incentive effects of power. *American Economic Review*, *103*(4), 1325–59.

Fershtman, C., & Gneezy, U. (2001). Strategic delegation: An experiment. *The RAND Journal of Economics*, *32*(2), 352-368.

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

Garfagnini, U., Ottaviani, M., & Sørensen, P. N. (2014). Accept or reject? an organizational perspective. *International Journal of Industrial Organization*, *34*, 66-74.

Gneezy, U. (2005). Deception: The role of consequences. *American Economic Review*, *95*, 384-394.

Gneezy, U., Rockenbach, B., & Serra-Garcia, M. (2013). Measuring lying aversion. *Journal of Economic Behavior and Organization*, *93*, 293-300.

Hamman, J. R., Loewenstein, G., & Weber, R. A. (2010). Self-interest through delegation: An additional rationale for the principal-agent relationship. *American Economic Review*, *100*(4), 1826–46.

Hamman, J. R., & Martínez-Carrasco, M. A. (2023). Managing uncertainty: An experiment on delegation and team selection. *Organization Science*, *34*(6), 2272-2295.

Holmstrom, B. (1984). On the theory of delegation. In M. Boyer & R. Kihlstrom (Eds.), *Bayesian models in economic theory.* North-Holland.

Jensen, M., & Meckling, W. (1992). Specific and general knowledge and organizational structure. In L. Werin & H. Wijkander (Eds.), *Contract economics.* Blackwell Publishers Oxford.

Kartik, N. (2009). Strategic communication with lying costs. *The Review of Economic Studies*, *76*(4), 1359-1395.

Lai, E. K., & Lim, W. (2012). Authority and communication in the laboratory. *Games and Economic*

Behavior, *74*(2), 541-560.

Lundquist, T., Ellingsen, T., Gribbe, E., & Johannesson, M. (2009). The aversion to lying. *Journal of Economic Behavior and Organization*, *70*, 81-92.

Osborne, M., & Rubinstein, A. (1994). *A course in game theory.* The MIT Press.

Ottaviani, M. (2000). *The economics of advice.* mimeo.

Owens, D., Grossman, Z., & Fackler, R. (2014). The control premium: A preference for payoff autonomy. *American Economic Journal: Microeconomics*, *6*(4), 138–61.

Özer, O., Subramanian, U., & Wang, Y. (2018). Information sharing, advice provision, or delegation. what leads to higher trust and trustworthiness? *Management Science*, *64*(1), 474–493.

Pitchik, C., & Schotter, A. (1987). Honesty in a model of strategic information transmission. *The American Economic Review*, *77*(5), 1032-1036.

Qin, X., Wang, S., & Wu, M. Z. (2024). Is it what you say or how you say it? *Experimental Economics*, *27*(4), 874–921. doi: 10.1007/s10683-024-09830-4

Rush, A., Smirnov, V., & Wait, A. (2010). Communication breakdown: Consultation or delegation from an expert with uncertain bias. *The BE Journal of Theoretical Economics*, *10*(1).

Serra-Garcia, M., Van Damme, E., & Potters, J. (2011). Hiding an inconvenient truth: Lies and vagueness. *Games and Economic Behavior*, *73*, 244-261.

Sloof, R., & von Siemens, F. A. (2021). Effective leadership and the allocation and exercise of power in organizations. *The Leadership Quarterly*, *32*(4), 101277. doi: 10.1016/j.leaqua.2018.12.001

Sobel, J. (2013). Giving and receiving advice. In D. Acemoglu, M. Arellano, & E. Dekel (Eds.), *Advances in economics and econometrics: Tenth world congress* (p. 305–341). Cambridge University Press.

Appendices

A Proofs of Proposition 1, 2, and 3

Proof of Proposition 1

We first show that Del-(3), Del-(2) and Del-(1) are the three relevant options to consider. Assuming that the worker always implements his most preferred project, it is never optimal for the manager to set $\Delta = \{B\}$, as this choice is dominated by $\Delta = \{A, B\}$. Clearly, $\Delta = \{C\}$ can never be optimal either, because project C is dominated by project B in expected payoff terms. Finally, $\Delta = \{A, C\}$ and $\Delta = \{B, C\}$ make no sense because full delegation would always yield the manager strictly more.

The optimal level of delegation thus follows from comparing the manager's expected payoffs in the three relevant cases. Under Del-(3) the manager earns $\frac{q}{3}(80 + 100 + 100) + (1 - q)10 = 10 + \frac{250}{3}q \equiv \Pi_{Del-(3)}$, under Del-(2) he gets $\frac{q}{3}(80 + 100 + 10) + (1 - q)80 = 80 - \frac{50}{3}q \equiv \Pi_{Del-(2)}$ and under Del-(1) he earns $\frac{q}{3}(80) + (1 - q)120 = 120 - \frac{280}{3}q \equiv \Pi_{Del-(1)}$. Comparing these payoffs, the thresholds for q are obtained. ■

Proof of Proposition 2.

Throughout the proof we define as 'equilibrium message' a message that is sent with positive probability in equilibrium. Moreover, we consider two equilibrium messages as being 'different' iff they induce a different probability distribution over project choices. The general structure of the proof is to first show that there is always an equilibrium message m_1 inducing project A for sure. Subsequently it is shown that in equilibria with two different equilibrium messages, types $t = 1$ and $t = 2$ necessarily fully separate, and if three (or more) different equilibrium messages are sent, types $t = 1$, $t = 2$ and $t = 3$ always fully separate.

We first show that in a neologism proof equilibrium, there necessarily exists an equilibrium message m_1 that induces project A for sure. Suppose not. Then type $t = 1$ could profitably send the neologism "Choose A, because my type belongs to 1". None of the other types would then have an incentive to mimic (as project A gives them their lowest possible payoff), while $t = 1$ strictly gains. Therefore, if only one message m_1 is sent in equilibrium (pooling), it must induce A. Types $t = 2$ through $t = 4$ then would strictly gain if they could induce the manager to choose B with positive probability instead; the neologism "Choose B, because my type belongs to $2, 3, 4$" is credible iff $q \geq \frac{12}{23}$. In a similar vein, the neologism "Choose C, because my type belongs to $3, 4$" is credible iff $q \geq \frac{33}{43}$. Hence pooling on m_1 is neologism proof iff $q < \frac{12}{23}$. This gives part 2(c).

Next consider equilibria with two different equilibrium messages, m_1 and m_2. Let α_{ia} be the probability that the manager chooses project $a \in \{A, B, C\}$ after message m_i. From the above,

$(\alpha_{1A}, \alpha_{1B}, \alpha_{1C}) = (1,0,0)$ necessarily. First, suppose $\alpha_{2B} > 0$. Then types $t = 2$ through $t = 4$ all strictly prefer, and thus choose, m_2 over m_1. After message m_2 the manager then prefers A if $q < \frac{12}{23}$ and B otherwise. The former case corresponds to the pooling outcome Com-(1) of part 2(c). The latter case corresponds to the Com-(2) outcome of part 2(b). The upper bound $q < \frac{33}{43}$ there follows from the neologism "Choose C, because my type belongs to $3,4$" being credible iff $q \geq \frac{33}{43}$. Second, suppose $\alpha_{2B} = 0$. For m_1 and m_2 to be different, then necessarily $\alpha_{2C} > 0$. This implies in turn that $t = 3$ and $t = 4$ choose m_2 for sure. If $\alpha_{2C} = 1$ type $t = 2$ could profitably send the neologism "Choose B, because my type belongs to 2". Hence necessarily $\alpha_{2C} < 1$, i.e. the manager strictly mixes between A and C. With $\beta_3(m_2) = \Pr(t = 3|m_2)$ and $\beta_4(m_2) = \Pr(t = 4|m_2)$ denoting the posterior beliefs, this requires $120\beta_4(m_2) = 100\beta_3(m_2) + 10\beta_4(m_2)$, i.e. $\frac{\beta_4(m_2)}{\beta_3(m_2)} = \frac{10}{11}$. From Bayes' rule, together with $t = 3$ and $t = 4$ choosing m_2 for sure, this in turn requires $\frac{(1-q)}{\frac{1}{3}q} = \frac{10}{11}$, i.e. $q = \frac{33}{43}$. But for this value of q either the neologism "Choose B, because my type belongs to 2" (if $\alpha_{2C} > \frac{7}{11}$), or the neologism "Choose B, because my type belongs to $2,4$" (if $\frac{1}{10} \leq \alpha_{2C} \leq \frac{7}{11}$), or the neologism "Choose B, because my type belongs to $2,3,4$" (if $\alpha_{2C} < \frac{1}{10}$) is credible.

Next consider equilibria with three different equilibrium messages. Let the three equilibrium messages be: m_1, m_2 and m_3, with $(\alpha_{1A}, \alpha_{1B}, \alpha_{1C}) = (1,0,0) \neq (\alpha_{2A}, \alpha_{2B}, \alpha_{2C}) \neq (\alpha_{3A}, \alpha_{3B}, \alpha_{3C})$. (Labeling of messages is such that $\alpha_{2B} \geq \alpha_{3B}$). This necessarily requires $\alpha_{2B} > \alpha_{3B}$ and $\alpha_{2C} < \alpha_{3C}$; if not, for all types $t = 2$ to $t = 4$ either m_2 or m_3 would be a dominated choice and not chosen in equilibrium. With $\alpha_{2B} > \alpha_{3B}$ type $t = 2$ will never choose message m_3. Let $\beta_j(m_3) = \Pr(t = j|m_3)$ again denote posterior beliefs. It then follows that $\beta_2(m_3) = 0$ and $\beta_3(m_3) + \beta_4(m_3) = 1$. Choosing project B after m_3 then gives the manager $80 - 70\beta_3(m_3)$ in expected payoffs, while project A yields him $120 - 120\beta_3(m_3)$. Project B yields more iff $\beta_3(m_3) > \frac{4}{5}$. But for these values of $\beta_3(m_3)$ project C is better for the manager (yielding $10 + 90\beta_3(m_3)$) than project B. Hence $\alpha_{3B} = 0$ necessarily. Next we will show that type $t = 3$ strictly prefers m_3. We will show this by contradiction. First we will show that if type $t = 3$ is indifferent between m_2 and m_3, then $t = 4$ strictly prefers m_2. Indifference of $t = 3$ implies $\alpha_{2B} = 10(\alpha_{3C} - \alpha_{2C})$. Choosing m_2 then yields type $t = 4$ an expected payoff of $10 + 70\alpha_{2B} + 110\alpha_{2C} = 10 + 700\alpha_{3C} - 590\alpha_{2C}$ while m_3 would yield an expected payoff of $10 + 110\alpha_{3C}$. For $\alpha_{3C} > \alpha_{2C}$ choosing m_2 yields type $t = 4$) more than choosing m_3. But in that case $\beta_3(m_3) = 1$ and $\alpha_{3C} = 1$. The latter is inconsistent with type $t = 3$ being indifferent between m_2 and m_3 as m_3 would yield type t_3 more when $\alpha_{3C} = 1 > \alpha_{2C}$. If type $t = 3$ strictly prefers m_2 over m_3, then $\beta_3(m_3) = 0$ and thus necessarily $\beta_4(m_3) = 1$; this would result in $\alpha_{3C} = 0$, contradicting $\alpha_{2C} < \alpha_{3C}$. Hence type $t = 3$ necessarily strictly prefers m_3. This in turn implies $\alpha_{2C} = 0$.

From the above it follows that, if three different equilibrium messages are sent in (a neologism proof) equilibrium, types $t = 1$, $t = 2$ and $t = 3$ fully separate by choosing m_1, m_2 and m_3, respectively.

Type $t = 4$ has no incentive to choose m_1 or yet another message m_4 that would separate him out, hence this type necessarily chooses between m_2 and m_3. If he does not choose m_3, $\beta_3(m_3) = 1$ and thus $\alpha_{3C} = 1$. But then $t = 4$ would prefer m_3 over m_2.

Therefore, type $t = 4$ sends m_3 for sure, or mixes between m_2 and m_3. Let us first consider the case where type $t = 4$ sends m_3 for sure. Suppose the manager observes m_3, then project C yields the manager $100\beta_3(m_3) + 10(1 - \beta_3(m_3))$ while project A yields $120(1 - \beta_3(m_3))$ (project B always gives less than the best of these two). For $\beta_3(m_3) = \frac{\frac{1}{3}q}{\frac{1}{3}q+(1-q)} \geq \frac{11}{21}$, it holds that the manager prefers project C over project A when m_3, hence $\alpha_{3C} = 1$. This holds when $q \geq \frac{33}{43}$ This gives part 2(a).

Next suppose that type t_4 mixes between m_2 and m_3 with probabilities τ_2 and τ_3 respectively. This requires mixing on the side of the manager. First, consider the manager mixing between action A and action C when m_3. This requires $\alpha_{3C} < 1$ and thus using Bayes' rule $\beta_3(m_3) = \frac{\frac{1}{3}q}{\frac{1}{3}q+\tau_3(1-q)} = \frac{11}{21}$. This reduces to $\tau_3 = \frac{10q}{33(1-q)}$. This in turn requires $q \leq \frac{33}{43}$ to secure $\tau_3 \leq 1$. Note that $\tau_3 = \frac{10q}{33(1-q)}$ implies $\tau_2 = \frac{33-43q}{33(1-q)}$. When m_2 is observed, the manager strictly prefers B over A if $\beta_2(m_2) > \frac{2}{7}$ implying $\alpha_{2B} = 1$. Choosing project B after m_2 yields the manager $80 + 20\beta_2(m_2)$ while choosing project A after m_2 yields the manager $120 - 120\beta_2(m_2)$. Using Bayes' rule $\beta_2(m_2) = \frac{\frac{1}{3}q}{\frac{1}{3}q+\tau_2(1-q)} = \frac{11q}{33-32q} > \frac{2}{7}$, resulting in $q > \frac{66}{141}$. In that case, for $t = 4$ to be willing to mix it is necessarily required that $10 + \alpha_{3C} \cdot 110 = 80$, i.e. $\alpha_{3C} = \frac{7}{11}$. This yields part 2(c). For $q < \frac{66}{141}$ the manager prefers project A after m_2, but then type t_4 strictly prefers to send m_3, and not to mix between m_2 and m_3.

Second, we will show that no mixed strategy equilibrium exists where type t_4 mixes between m_2 and m_3 and the manager mixes between action A and action B when m_2. Mixing between A and B when m_2 requires $\alpha_{2B} < 1$ and thus using Bayes' rule $\beta_2(m_2) = \frac{\frac{1}{3}q}{\frac{1}{3}q+\tau_2(1-q)} = \frac{2}{7}$. This reduces to $\tau_2 = \frac{5q}{6(1-q)}$. Note that $\tau_2 = \frac{5q}{6(1-q)}$ implies $\tau_3 = \frac{6-11q}{6(1-q)}$. When m_3 is observed, the manager strictly prefers C over A if $\beta_3(m_3) > \frac{11}{21}$ implying $\alpha_{3C} = 1$. Using Bayes' rule $\beta_3(m_3) = \frac{\frac{1}{3}q}{\frac{1}{3}q+\tau_3(1-q)} = \frac{2q}{6-9q} > \frac{11}{21}$, resulting in $q > \frac{66}{141}$. But if $\alpha_{3C} = 1$, then type t_4 strictly prefers to send m_3, and not to mix between m_2 and m_3. For $q < \frac{66}{141}$ the manager strictly prefers project A over C when m_3. But in that case type t_4 strictly prefers to send m_2, and not to mix between m_2 and m_3. Hence, no equilibrium exists where t_4 mixes between m_2 and m_3 and the manager mixes between action A and action B when m_2. ■

Proof of Proposition 3.

Proposition 3 compares delegation and communication in terms of expected payoffs for the manager.

The payoffs of under delegation are as follows. Under Del-(3) the manager earns $\frac{q}{3}(80+100+100)+(1-q)10 = 10+\frac{250}{3}q \equiv \Pi_{Del-(3)}$, under Del-(2) he gets $\frac{q}{3}(80+100+10)+(1-q)80 = 80-\frac{50}{3}q \equiv \Pi_{Del-(2)}$ and under Del-(1) he earns $\frac{q}{3}(80) + (1-q)120 = 120 - \frac{280}{3}q \equiv \Pi_{Del-(1)}$.

The payoffs under communication are as follows. Under Com-(3) the manager earns $\frac{q}{3}(80 + 100 + 100) + (1-q)10 = 10 + \frac{250}{3}q \equiv \Pi_{Com-(3)}$, under Com-(2) he gets $\frac{q}{3}(80 + 100 + 10) + (1-q)80 =$

$80 - \frac{50}{3}q \equiv \Pi_{Com-(2)}$ and under Com-(1) he earns $\frac{q}{3}(80) + (1-q)120 = 120 - \frac{280}{3}q \equiv \Pi_{Com-(1)}$. Finally, under the mixed equilibrium Com-(3-mixed) the manager earns $80 - \frac{260}{33}q \equiv \Pi_{Com-(3-mixed)}$.

Note that $\Pi_{Del-(3)} = \Pi_{Com-(3)}$, $\Pi_{Del-(2)} = \Pi_{Com-(2)}$, and $\Pi_{Del-(1)} = \Pi_{Com-(1)}$. The mixed equilibrium Com-(3-mixed) exceeds $\Pi_{Del-(3)}$ whenever $q \leq \frac{33}{43}$, it always exceeds $\Pi_{Del-(2)}$, and it exceeds $\Pi_{Del-(1)}$ whenever $q \geq \frac{66}{141}$. ■

B Predictions under lying aversion and credulity

In this Appendix we derive the equilibrium predictions in the presence of some (at least) modest fraction of credulous managers γ who always believe and act upon the message received, as well as some fraction λ of lying averse workers that always report the true state. In particular, we verify the extent to which the equilibrium outcomes listed in Proposition 2 as description of the strategic types' equilibrium behavior (in the absence of non-strategic types), is robust to having fractions of such credulous and lying averse types in the population.

Proposition 4 *Suppose that a fraction $\lambda \in (0,1)$ of workers is lying averse and a fraction $\gamma > \frac{1}{10}$ of managers is credulous. Outcomes Com-(1) and Com-(2) in Proposition 2 with No and Limited Communication then no longer describe equilibrium behavior of the strategic workers and strategic managers. Strategic workers thus always send a truthful message in the congruent states (i.e. message m_t in states $t = 1, 2, 3$) while the strategic manager always chooses A in response to m_1 and m_4. The following Communication equilibria exist:*

(a) *For $q > \frac{33(1-\lambda)}{10+33(1-\lambda)}$, Com-(3): In the conflict state $t = 4$ the strategic worker sends message m_3. The strategic manager chooses B in response to m_2 and C in response to m_3;*

(b) *For $q < \frac{33(1-\lambda)}{10+33(1-\lambda)}$ and $\gamma > \frac{7}{11}$, Com-(3A): In the conflict state $t = 4$ the strategic worker sends message m_3. The strategic manager chooses B in response to m_2 and A in response to m_3.*

(c) *For $\frac{66(1-\lambda)}{141-66\lambda} < q < \frac{33(1-\lambda)}{10+33(1-\lambda)}$ and $\gamma < \frac{7}{11}$, Com-(3-mixed): In the conflict state $t = 4$ the strategic worker sends message m_2 with probability $\tau_2 = \frac{33(1-\lambda(1-q))-43q}{33(1-\lambda)(1-q)}$ and m_3 with probability $\tau_3 = \frac{10q}{33(1-\lambda)(1-q)}$. The strategic manager chooses B after m_2 and after m_3 she chooses A with probability $\alpha_{3A} = \frac{4}{11(1-\gamma)}$ and C with probability $\alpha_{3C} = \frac{7-11\gamma}{11(1-\gamma)}$.*

(d) *For $q < \frac{66(1-\lambda)}{141-66\lambda}$ and $\gamma < \frac{7}{11}$, Com-(3A-mixed): In the conflict state $t = 4$ the strategic worker sends message m_2 with probability $\tau_2 = \frac{5q}{6(1-\lambda)(1-q)}$ and m_3 with probability $\tau_3 = \frac{6(1-\lambda)(1-q)-5q}{6(1-\lambda)(1-q)}$. The strategic manager chooses A with probability $\alpha_{2A} = \frac{7-11\gamma}{7(1-\gamma)}$ and B with probability $\alpha_{2B} = \frac{4\gamma}{7(1-\gamma)}$ after m_2, and after m_3 she chooses A.*

As noted in the main text, there are two major changes relative to the standard equilibrium predictions presented in Proposition 2. First, all workers – strategic and lying averse alike – always tell the truth in all congruent states, so Com-(1) and Com-(2) no longer exist. Second, the equilibrium predictions shift to project A becoming a best response for strategic managers more often. To illustrate the latter, for a large fraction of credulous managers $\gamma > \frac{7}{11}$, strategic workers in state $t = 4$ prefer to always send m_3. Although this induces strategic managers to choose A in response to m_3, workers are very likely to encounter a credulous manager instead who chooses C in response to m_3. This leads to an additional equilibrium Com-(3A) that is equivalent to Com-(3) in terms of worker behavior, but differs in the strategic managers' response to m_3.[42] In a similar vein, for lower levels of interest alignment now an additional mixed strategy equilibrium Com-(3A-mixed) exists in which a strategic manager now chooses project A for sure after m_3 (rather than mixes between A and C as in Com-(3-mixed)), and mixes between projects A and B in response to m_2 (rather than project B for sure).[43]

Proof of Proposition 4.
By definition, the fraction λ of lying averse workers always sends message m_1 if t_1, m_2 if t_2, and m_3 if t_3 and m_4 if t_4. Moreover, the fraction γ of credulous managers always chooses A if m_1, B if m_2, C if m_3, and A if m_4. The remainder of the proof thus focuses on the behavior of strategic workers and strategic managers.

We first show that the Com-(1) equilibrium in Proposition 2 is not robust to having a fraction of $\gamma > 0$ of credulous managers. Suppose strategic workers always send message m_1, leading to project A in response. Then by deviating to message m_2, a strategic worker of type $t = 2$ can obtain at least $100\gamma > 0$ in expected payoff terms, because credulous managers will take m_2 at face value and choose B in response. In a similar vein, in the Com-(2) equilibrium of Proposition 2 the strategic worker of type $t = 3$ obtains an expected payoff of 10. Deviating to message m_3 then yields at least 100γ and is thus certainly profitable whenever $\gamma > \frac{1}{10}$. This leaves the equilibria Com-(3) and Com-(3-mixed) from Proposition 2 in which the strategic worker always sends a truthful message in the three congruent states and the strategic manager always chooses A in response to m_1. These two equilibria differ in whether the strategic worker uses a pure or a mixed strategy in conflict state $t = 4$.

Consider Com-(3) in which the strategic worker uses a pure strategy first. If the strategic worker always chooses m_3 whenever $t = 3$ or $t = 4$ (and never otherwise), the rational posterior belief of a strategic manager that the state equals $t = 3$ (rather than $t = 4$) upon observing m_3 equals $\beta_3(m_3) = \frac{\frac{q}{3}}{\frac{q}{3}+(1-\lambda)(1-q)}$. Choosing A in response to m_3 then yields $120-120\beta_3(m_3)$, B gives $80-70\beta_3(m_3)$, while choosing C gives $10 + 90\beta_3(m_3)$. Project B always gives less than the best of A and C does. Project

[42] For the standard Crawford and Sobel framework, Ottaviani (2000) shows that there is a fully revealing equilibrium as soon as there is a positive fraction of credulous (or 'naive' in his wording) managers.

[43] This additional equilibrium collapses to the Com-(1) outcome for $\gamma = 0$.

C is preferred over A whenever $\beta_3(m_3) > \frac{11}{21}$. Rewriting gives $q > \frac{33(1-\lambda)}{10+33(1-\lambda)}$ and thus part (a). If this condition is not satisfied, the strategic manager will choose A in response to m_3. If so, for the strategic worker to strictly prefer m_3 over m_2 in conflict state $t = 4$, it must hold that $10+110\gamma > 80$, i.e. $\gamma > \frac{7}{11}$. This gives the alternative Com-(3A) equilibrium of part (b).

Next, suppose that the type $t = 4$ strategic worker mixes between m_2 and m_3 with probabilities τ_2 and τ_3, respectively. This requires mixing on the side of the manager. First, consider the manager mixing between actions A action C when m_3. The managers' indifference between these two actions requires $\beta_3(m_3) = \frac{\frac{q}{3}}{\frac{q}{3}+(1-\lambda)(\tau_3(1-q))} = \frac{11}{21}$ per the above and Bayes' rule. This reduces to $\tau_3 = \frac{10q}{33(1-\lambda)(1-q)}$. This in turn requires $q < \frac{33(1-\lambda)}{10+33(1-\lambda)}$ to secure $\tau_3 < 1$. Furthermore, $\tau_3 = \frac{10q}{33(1-\lambda)(1-q)}$ implies $\tau_2 = \frac{33-43q-33\lambda(1-q)}{33(1-\lambda)(1-q)}$ and thus $\beta_2(m_2) = \frac{\frac{q}{3}}{\frac{q}{3}+(1-\lambda)\tau_2(1-q)} = \frac{11q}{q+33(1-\lambda)(1-q)}$ by Bayes' rule. If $\beta_2(m_2) > \frac{2}{7}$ the manager strictly prefers B over A after m_2 implying $\alpha_{2B} = 1$. This translates to $q > \frac{66(1-\lambda)}{141-66\lambda}$. In that case, for $t = 4$ to be willing to mix, it is necessarily required that $10 + 110\gamma + (1-\gamma)110\alpha_{3C} = 80$, i.e. $\alpha_{3C} = \frac{7-11\gamma}{11(1-\gamma)}$ and $\alpha_{3A} = 1 - \alpha_{3C} = \frac{4}{11(1-\gamma)}$. Note that $\alpha_{3C} > 0$ requires $\gamma < \frac{7}{11}$. This gives part 4(c).

Second, we consider whether a mixed strategy equilibrium exists where type $t = 4$ mixes between m_2 and m_3 and the manager mixes between actions A and B in response to m_2. The latter requires $\beta_2(m_2) = \frac{\frac{q}{3}}{\frac{q}{3}+(1-\lambda)\tau_2(1-q)} = \frac{2}{7}$ per the above and Bayes' rule. This reduces to $\tau_2 = \frac{5q}{6(1-\lambda)(1-q)}$. This in turn requires $q < \frac{6(1-\lambda)}{11-6\lambda}$ to secure $\tau_2 < 1$. Note that $\tau_2 = \frac{5q}{6(1-\lambda)(1-q)}$ implies $\tau_3 = \frac{6(1-\lambda)(1-q)-5q}{6(1-\lambda)(1-q)}$. When m_3 is observed, the manager strictly prefers C over A if $\beta_3(m_3) > \frac{11}{21}$, implying $\alpha_{3C} = 1$. Using Bayes' rule $\beta_3(m_3) = \frac{\frac{q}{3}}{\frac{q}{3}+(1-\lambda)\tau_3(1-q))} = \frac{2q}{2q+6(1-\lambda)(1-q)-5q} > \frac{11}{21}$ if $q > \frac{66(1-\lambda)}{141-66\lambda}$. But in case $\alpha_{3C} = 1$, the strategic worker of type $t = 4$ strictly prefers to send m_3 over mixing between m_2 and m_3. For $q < \frac{66(1-\lambda)}{141-66\lambda}$ the strategic manager strictly prefers project A over C when m_3. Then choosing m_3 yields type $t = 4$ an expected payoff of $120\gamma + 10(1-\gamma)$ while choosing m_2 yields an expected payoff of $80\gamma+(10+70\alpha_{2B})(1-\gamma)$. Type $t = 4$ is willing to mix iff $120\gamma+10(1-\gamma) = 80\gamma+(10+70\alpha_{2B})(1-\gamma)$, i.e. $\alpha_{2B} = \frac{4\gamma}{7(1-\gamma)}$ and $\alpha_{2A} = 1 - \alpha_{2B} = \frac{7-11\gamma}{7(1-\gamma)}$. Note that $\alpha_{2A} > 0$ requires $\gamma < \frac{7}{11}$. (Moreover, $q < \frac{66(1-\lambda)}{141-66\lambda}$ implies $q < \frac{6(1-\lambda)}{11-6\lambda}$.) This gives part 4(d). ■

C Improved organizational decision-making

In this Appendix, we examine whether centralization and delegation enable managers to leverage workers' private information in a way that is profitable for the organization, and how this depends on the degree of interest alignment. We use the no-information case as a benchmark. In the absence of any additional information about the state of the world, managers should always choose project A in our experiment, since $q < \frac{4}{5}$ in all treatments. This yields an expected payoff of $120 - \frac{280}{3}q$ for the manager, with project A being the manager's ex post optimal choice in a fraction $1 - \frac{2}{3}q$ of cases. By construction, these benchmarks decrease with interest alignment q, because managers earn the highest payoff when interests conflict ($t = 4$) and project A is chosen. To cleanly isolate how the benefits of communication and delegation vary with interest alignment, we therefore measure improvements in organizational decision making relative to the no-information benchmark.

Theory (cf. Section 2) predicts that the benefits from improved organizational decision making increase with the degree of interest alignment. The intuition is straightforward. In terms of the implemented project choice, the no-information benchmark coincides with the Com-(1) and Del-(1) equilibrium benchmarks. Hence, for low levels of interest alignment, both centralization and delegation are predicted to perform no better than the no-information benchmark. For higher levels of q, however, they are predicted to perform increasingly better relative to this benchmark, because the gains associated with the alternative communication and delegation equilibria, relative to Com-(1) and Del-(1), are themselves increasing in q.

Table C.1: Improved organizational decision making (vs. no-information benchmark)

	$q = 0.75$	$q = 0.6$	$q = 0.4$	p-value
Managers' payoffs				
Δ payoffs Centralization	23.0	9.6	-3.5	< 0.01
Δ payoffs Delegation	21.2	0.9	-9.5	< 0.01
p-value C vs. D	0.59	< 0.01	0.11	
Managers' preferred project				
Δ % correct choice Centralization	16%	-1%	-20%	< 0.01
Δ % correct choice Delegation	15%	-17%	-28%	< 0.01
p-value C vs. D	0.99	< 0.01	0.05	

Notes: The table reports the improvement (denoted as Δ) in terms of managers' payoffs and preferred project relative to the counterfactual, no-information benchmark case in which always project A is chosen. The reported p-values in the rightmost column follows from a Cuzick test for trend. Those reported in the rows labeled 'p-values C. vs. D' compare the improvements under Centralization and Delegation by means of Wilcoxon signed-rank tests for matched pairs. All statistical tests are performed on the distributions of matching group averages (including data from both waves whenever possible). All data is taken from parts I and II where organizational design is exogenously imposed.

The two measures in Table C.1 capture related but distinct aspects of decision quality. Managers' payoffs are the primary payoff-relevant criterion. The preferred-project measure is useful as a simple indicator of whether the ex post optimal project is implemented, but it treats all deviations from that project equally. In the payoff matrix of Table 1, however, mistakes differ in severity. For example, in the conflict state $t = 4$ choosing the compromise project B rather than A is less costly for the manager than choosing C, even though both are counted as incorrect. Conversely, a change in behavior may improve expected payoffs without increasing the frequency with which the manager's most preferred project is chosen. We therefore interpret the preferred-project measure as complementary to, rather than a substitute for, the payoff measure.

The data indicate for both measures that managers perform increasingly better relative to the no-information benchmark as interests get better aligned. Specifically, managers' relative payoff gains increase with the level of interest alignment according to a Cuzick trend test ($p < 0.01$), both under centralization and under delegation.[44] We also find that in case of medium alignment, managers' payoff gains are significantly higher under centralization than under delegation ($p < 0.01$, Wilcoxon signed-rank test),[45] whereas no significant differences between organizational forms are found under strong or weak alignment. Results for the second, managers' preferred project choice measure are qualitatively similar.[46] The only difference is that, in case of weak alignment, the extent to which the manager's preferred project is implemented relative to the no-information benchmark is now also significantly higher under centralization than under delegation ($p = 0.05$; Wilcoxon signed-rank test). Overall we conclude that, as interests become better aligned, decision quality improves in both dimensions, by improving expected payoffs and the frequency with which the manager's most preferred project is chosen.

[44]Pairwise comparisons conducted separately for each wave display a similar pattern. In wave 1, the manager's payoff relative to the no-information benchmark decreases significantly ($p < 0.01$, ranksum test) when going from strong to medium alignment, both under centralization (23.0 versus 10.4) and under delegation (21.2 versus 1.4). In wave 2, payoff gains decrease significantly ($p < 0.01$, ranksum test) when going from medium to weak alignment, both under centralization (8.9 versus -3.5) and under delegation (0.5 versus -9.5).

[45]This also holds when considering wave 1 and wave 2 separately.

[46]Pairwise comparisons conducted separately for each wave display a similar pattern. In wave 1, the extent to which the manager's preferred project is implemented, relative to the no-information benchmark, decreases significantly ($p < 0.01$, ranksum test) when going from strong to medium alignment, both under centralization (16% versus -1%) and under delegation (15% versus -17%). In wave 2, the gain in % of correct project choices decreases significantly ($p < 0.01$, ranksum test) when going from medium to weak alignment, both under centralization (-1% versus -20%) and under delegation (-17% versus -28%). Additionally, under medium alignment, the %-gain in correct choices is significantly higher under centralization than under delegation ($p < 0.01$ for both waves 1 and 2; Wilcoxon signed-rank test).

D Best response analysis

D.1 Best response of workers' messages

Table D.1: Workers' messages and best responses by treatment and state

	State 1		State 2		State 3		State 4	
	Observed	BR	Observed	BR	Observed	BR	Observed	BR
Strong alignment (q = 0.75)								
M1A	98	100 [98]	0	0 [0]	1	0 [0]	1	0 [0]
M2B	2	0 [0]	99	100 [99]	2	0 [0]	43	69 [32]
M3C	0	0 [0]	0	0 [0]	93	100 [93]	54	31 [17]
M4N	0	n.a.	1	n.a.	4	n.a.	2	n.a.
Medium alignment (q = 0.6)								
M1A	96	100 [96]	0	0 [0]	0	0 [0]	1	0 [0]
M2B	0	0 [0]	99	100 [99]	4	8 [0]	61	76 [44]
M3C	1	0 [0]	0	0 [0]	92	92 [84]	31	24 [6]
M4N	3	n.a.	1	n.a.	4	n.a.	7	n.a.
Weak alignment (q = 0.4)								
M1A	96	100 [96]	0	0 [0]	0	0 [0]	0	0 [0]
M2B	2	0 [0]	98	100 [98]	0	0 [0]	62	58 [34]
M3C	1	0 [0]	0	0 [0]	94	100 [94]	26	42 [11]
M4N	2	n.a.	2	n.a.	6	n.a.	12	n.a.

Note: Percentages in table are based on matching group averages. The columns labeled Observed reports the observed percentage of cases in which the message in question was sent. The columns labeled BR report the percentage of cases in which the message in question would be the worker's best response. Numbers within brackets report, in percentage points, the overlap between actually observed messages and best-response messages, using the same treatment–state denominator as the corresponding Observed and BR entries. n.a. means that best responses are not available because of insufficient actual occurrences of messages M4N.

One remark applies regarding the identification of best responses in the above table. In a number of instances (i.e. matching groups) two different messages induce the manager to make the same choice. For instance, after both M1A and M4N some managers choose project A. In that case there sometimes is no unique best response message when the state is 1. For ease of presentation we do not label these cases separately, but include them all in category M1A when the state is 1, as M1A is always one of these multiple best responses. More precisely, under strong, medium, and weak alignment, the category M1A in the BR column column is a combined fraction: 77%, 72%, and 58% respectively have M1A as unique best response, while 23%, 28%, and 42% respectively have M1A and M4N as best response. Similarly, after both M2B and M4N some managers choose project B. In that case there may not be a unique best response when the state is 2. Also these are not labelled separately in the table, but included in category M2B when the state is 2. The category M2B in the BR column is thus also a combined fraction; 77%, 92%, and 92% respectively under strong, medium and weak interest alignment have M2B as unique best response, while 23%, 8%, and 8% respectively have both M2B and M4N as best response.

D.2 Best response analysis of managers' responses to messages

Table D.2: Managers' responses to received messages and best responses by treatment

	Project A		Project B		Project C	
	Observed	BR	Observed	BR	Observed	BR
Strong alignment (q = 0.75)						
M1A	99	100 [99]	0	0 [0]	1	0 [0]
M2B	2	0 [0]	98	100 [98]	0	0 [0]
M3C	42	0 [0]	9	0 [0]	50	100 [50]
M4N	44	n.a.	52	n.a.	4	n.a.
Medium alignment(q = 0.6)						
M1A	99	100 [99]	0	0 [0]	1	0 [0]
M2B	9	0 [0]	91	100 [91]	0	0 [0]
M3C	48	20 [9]	17	0 [0]	35	80 [28]
M4N	68	n.a.	32	n.a.	0	n.a.
Weak alignment (q = 0.4)						
M1A	98	100 [98]	2	0 [0]	0	0 [0]
M2B	38	58 [25]	61	42 [29]	0	0 [0]
M3C	60	83 [50]	17	0 [0]	23	17 [3]
M4N	83	n.a.	18	n.a.	0	n.a.

Note: Percentages in table are based on matching group averages. The columns labeled Observed reports the observed percentage of cases in which the project in question was chosen. The columns labeled BR report the percentage of cases in which the project in question would be the manager's best response. Numbers within brackets report, in percentage points, the overlap between actually chosen projects and best-response projects, using the same treatment–message denominator as the corresponding Observed and BR entries. n.a. means that best responses are not available because of insufficient actual occurrences of messages M4N.

E Factual statements versus recommendations

Table E.1: Workers' messages by treatment and state in wave 1

	Strong				**Medium**			
	State 1	State 2	State 3	State 4	State 1	State 2	State 3	State 4
M1	84	0	0	0	77	0	0	0
M2	0	84	2	30	0	82	5	50
M3	0	0	74	45	0	0	76	21
M4	0	1	1	1	0	0	1	2
MA	14	0	1	1	18	0	0	1
MB	2	16	0	13	0	17	1	16
MC	0	0	19	9	2	0	16	7
MN	0	0	3	2	3	1	2	2

Note: The table reflects the percentages based on matching group averages.

Table E.2: Worker's messages by state and treatment in wave 2

	Medium-S				**Weak-S**			
	State 1	State 2	State 3	State 4	State 1	State 2	State 3	State 4
M1	93	0	0	1	95	0	0	0
M2	0	99	1	43	3	96	0	60
M3	0	0	94	40	2	0	98	24
M4	3	1	1	13	0	4	0	7
MN	4	0	3	3	0	0	2	9
	Medium-R				**Weak-R**			
	State 1	State 2	State 3	State 4	State 1	State 2	State 3	State 4
MA	100	0	0	0	97	0	0	0
MB	0	100	6	71	0	100	0	65
MC	0	0	89	27	0	0	91	28
MN	0	0	5	3	3	0	9	7

Note: The table reflects the percentages based on matching group averages.

F Sample instructions [for sessions of first row in Table 2]

Instructions Experiment

General Information

Thank you for participating in this experiment. The amount of money you earn depends upon the decisions you and the other participants make. Your earnings are given in points. The experiment consists of four parts. Your overall earnings equal the sum of your points in each part. The conversion rate is 2.5 points for 1 eurocent, so 250 points corresponds with 1 euro. We will pay out your overall earnings in cash after you have completed the entire experiment and filled out a final questionnaire. We ensure that your final earnings remain confidential: no other participant from the experiment will learn your final earnings.
This (two-sided) sheet contains the instructions for part one of the experiment. Instructions for the next part follow after this part has been completed (and so on). Please do not communicate with other participants during the experiment. If you have a question, please raise your hand. The experimenter will then come to your table to answer your question in private.

Instructions part I

There are two types of participants: managers and workers. One half of the participants will be managers, and the remaining half will be workers. You will be randomly assigned one of these roles. Which role you have, you will learn at the start of this part. You will keep the same role in parts II and III. In part one 20 project implementation decisions have to be made. In every project implementation decision (period), manager and worker face three projects **(A,B,C)** that can be implemented. These projects differ in the points that they yield manager and worker upon implementation. The points belonging to a given project depend on the state of the world. There are four possible states **(1,2,3,4)**. The following table presents the points the different projects yield manager and worker in the different states:

	Project **A**		Project **B**		Project **C**	
	Manager	Worker	Manager	Worker	Manager	Worker
State **1**	80	80	0	0	0	0
State **2**	0	0	100	100	0	0
State **3**	0	0	10	10	100	100
State **4**	120	10	80	80	10	120

At the beginning of each period the computer determines the state. The four states are equally likely. That is, with probability 25% the state is 1, with probability 25% the state is 2, with probability 25% the state is 3, and with probability 25% the state is 4.

As explained above, Part I consists of 20 periods. In each period, one manager and one worker are randomly paired. You are never paired with the same other participant twice in a row. You cannot predict when you will be paired with the same other participant again.

At the beginning of each period the worker learns the state. The manager decides which project (either A, B, or C) will be implemented. Before the manager does so, the worker sends a message to the manager. The manager only learns the actual state at the end of the period, after the project has been implemented.

Sequence of Actions

The precise timing within each period is as follows. There are two phases.

Phase 1

The worker learns the state and sends one of the following eight messages to the manager:

The state is 1
The state is 2
The state is 3
The state is 4
I recommend project A
I recommend project B
I recommend project C
I make no recommendation

The set of available messages does not depend on the actual state; so irrespective of the actual state, the worker can always choose one of the above eight messages.

Phase 2

The manager observes the worker's message (but not the actual state) and decides which project to implement: either A, B, or C.

Upon completion of phase 2 both manager and worker are informed about the outcome in that period.

Payoffs

The number of points earned by manager and worker respectively are their points from the implemented project; see the table on the other side of the sheet.
Your overall payoff from this part is the sum of points earned in the 20 periods.

Instructions Part II

This part of the experiment also consists of 20 periods. As compared to part I the main difference is that now the worker takes the implementation decision in each period. Before s/he does so, the manager may restrict the set of projects that the worker can choose from.
As before, in each period one manager and one worker are randomly paired. You are never paired with the same other participant twice in a row. You cannot predict when you will be paired with the same other participant again.
Like in part I the worker learns the state (the four different states occur with the same probabilities as in part I). Before the worker decides which project to implement, the manager decides which projects the worker is allowed to choose for implementation. The manager only learns the state at the end of the period, after the project has been implemented.

Sequence of Actions

The precise timing within each period is as follows. There are two phases.

Phase 1

The manager selects the projects that the worker is allowed to choose for implementation.

Phase 2

The worker learns the state and the set of projects from which s/he is allowed to choose, and decides which of these projects to implement. (If the manager only allows one project, this becomes a trivial choice.)
Upon completion of phase 2 both manager and worker are informed about the outcome in that period.

Payoffs

The number of points earned by manager and worker respectively are their points from the implemented project. The same table as in part I applies.
Your overall payoff from this part is the sum of points earned in the 20 periods.

Instructions Part III

This part of the experiment combines the previous two parts and consists of 20 project implementation decisions (periods). At the beginning of each period the manager first decides whether s/he will take the implementation decision her- or himself or whether s/he delegates the implementation decision to the worker. The former corresponds with the decision structure of part I, the latter with the decision structure of part II.
As in the previous parts, in each period one manager and one worker are randomly paired. You are never paired with the same other participant twice in a row. You cannot predict when you will be paired with the same other participant again.

Sequence of Actions

The precise timing within each period is as follows. There are three phases.

Phase 0

The manager chooses between taking the project implementation decision her- or himself, **Decision Manager**, and delegation of this decision to the worker, **Decision Worker**.

Phases 1 and 2

The structure of phases 1 and 2 depends on the manager's choice in phase 0:

	Decision Manager (phases 1 and 2 are as in part I)	**Decision Worker** (phases 1 and 2 are as in part II)
Phase 1	The worker learns the state and sends a message(out of the eight messages possible) to the manager;	The manager selects the projects that the worker is allowed to choose for implementation
Phase 2	The manager observes the worker's message (but not the actual state) and decides which project to implement	The worker learns the state and the set of projects from which s/he is allowed to choose, and decides which of these projects to implement

Payoffs

The number of points earned by manager and worker respectively are their points from the implemented project. The same payoff table as in parts I and II applies. Your overall payoff from this part is the sum of points earned in the 20 periods.

Overview parts I and II

Before the start of the first period in part III, you will get a history overview of your own outcomes in parts I and II.

Instructions Part IV

The final part has a structure that differs from the previous parts. It consists of only one period. You are randomly paired with one other participant, taking either the role of participant A or participant B.

The computer will randomly assign one of the numbers 1, 2, 3, 4, 5, or 6 to each pair. Based on the assigned number, participant A sends a message about that number to Participant B. Participant B observes the message sent by Participant A, but not the number actually assigned by the computer, and decides whether or not to follow the message of Participant A.

Both participants A and B indicate their choices for each situation that may occur. In particular, for every possible number assigned to the pair, participant A has to formulate a message to participant B about the assigned number. The message does not have to contain the actual assigned number:

If the assigned number is	1	2	3	4	5	6
then your message to participant B is: The assigned number is	---	---	---	---	---	---

Before participant B receives the actual message of participant A, s/he has to decide for all possible messages whether or not s/he will follow the message:

If participant A sends the message	Then your decision is:
"The assigned number is 1"	to follow O O not to follow
"The assigned number is 2"	to follow O O not to follow
"The assigned number is 3"	to follow O O not to follow
"The assigned number is 4"	to follow O O not to follow
"The assigned number is 5"	to follow O O not to follow
"The assigned number is 6"	to follow O O not to follow

After A and B have made their decisions, the choices that result from the number actually assigned to the pair are carried out. That is, given the assigned number, the corresponding message of participant A is sent to participant B and the decision of participant B corresponding to this message is carried out.

Payoffs

Participant A receives 500 points plus 100 times the number sent in the message to Participant B. This means that if Participant A sends the message that the assigned number is 1, then her/his payoff equals 600 points; if Participant A sends the message that the number is 2, then her/his payoff equals 700 points; and so on.

The payoff of participant B depends on whether s/he decides to follow Participant A or not and on whether the message of Participant A contains the actual assigned number to the pair.

If Participant B decides to **follow**, then Participant B receives 500 points if the message of Participant A contains the actual assigned number to the pair. Otherwise Participant B receives zero. This means that if Participant B decides to follow and the actual assigned number to the group is 1 and the reported number in the message of Participant A is 1, then Participant B gets 500 points; however, if Participant B decides to follow and the actual assigned number to the group is 1 and the reported number in the message of Participant A is 2, then Participant B earns zero points; and so on.

If Participant B decides **not to follow**, then Participant B receives 150 points.