Truth Be Told

An Experimental Study of Communication and Centralization

Jordi Brandts
David J. Cooper

June 2018

Barcelona GSE Working Paper Series

Working Paper nº 1046

Truth Be Told
An Experimental Study of Communication and Centralization

by Jordi Brandts* and David J. Cooper**

* Instituto de Análisis Económico (CSIC) and Barcelona GSE
Email: jordi.brandts@iae.csic.es
Mail: Campus UAB, 08193 Bellaterra (Barcelona), Spain

** Florida State University and University of East Anglia
Email: djcooper@fsu.edu
Mail: 265 Bellamy Building, Florida State University, Tallahassee, FL 32306, USA

May 28, 2018

Abstract: We study the tradeoffs between centralized and decentralized management using a new experimental game, the decentralization game. Product types for two divisions are either chosen independently by the divisions (decentralization) or imposed by a central manager (centralization). Centralization makes it easier to coordinate the divisions' product types but more difficult to take advantage of the divisions' private information. We find that total surplus is highest when centralization is combined with free-form chat between the three players. This high performance occurs because divisions almost never lie about their private information, yielding unambiguous transmission of information from divisions to the central manager.

Keywords: Coordination, Experiments, Organizations, Asymmetric Information

JEL Classification Codes: C92, D23, J31, L23, M52

Acknowledgements: The authors thank the National Science Foundation (SES-0214310), the Spanish Ministry of Economy and Competitiveness (Grant: ECO2017-88130), the Severo Ochoa Program for Centers of Excellence in R&D (SEV2015-0563), and the *Generalitat de Catalunya* (Grant: 2017 SGR 1136). We thank Adrià Bronchal, Joe Ballard, Ellis Magee, and Lavinia Piemontese for valuable help as research assistants and seminar participants at Copenhagen Business School, Durham, East Anglia, Edinburgh, Florida State, Lausanne, Paderborn, Paris School of Economics, Stavanger, the London Experimental Workshop and the University of California – Santa Barbara for useful feedback.

1. Introduction: Should organizations be centrally managed or should units within the organization operate independently? This is an old question in economics that plays a central role in the economics of organizations. Consider, for example, a large firm that must choose product types to be sold in different locations. Should a central manager assign product types to each location or should managers pick? As another example, think of a small company choosing what type of software the workers will use (e.g., MS Word or LaTeX). Should the owner impose a choice on their employees or should the employees be allowed to choose? There are many other examples of this kind.

Answering such questions often involves a tradeoff between coordination and efficient use of information. In both of our examples there are advantages to subordinates coordinating their choices: Economies of scale imply that costs are lower for a large firm if its local divisions sell similar product types, and it is easier for workers to collaborate if they use the same software. Coordination is problematic because subordinates may have differing opinions about what option best fits their needs. Local divisions will want to choose product types that conform to local tastes. Workers will want software with which they are already familiar. If each subordinate is free to choose an option, they are unlikely to spontaneously coordinate on a single choice. Having the manager impose a choice solves this coordination problem.

The preceding suggests that centralization is the obvious choice, but suppose the manager knows less than her subordinates about the relative benefits of the various options. Local managers are likely better informed about demand conditions than a central manager at a remote location, and workers who actually use the software are probably better informed about its merits than the firm's owner. The manager can ask her subordinates for input, but they have incentives to exaggerate the benefits of their preferred option as a way of influencing the manager's choice. Much of the recent literature studying the tradeoffs between centralization and decentralization sees this tradeoff in terms of a comparison between the benefits of coordination and the costs of distorted information that accompany centralization.[1]

The purpose of this paper is to compare performance under decentralization and centralization in a demanding experimental environment that accentuates differences between the two organizational forms. Our primary findings are as follows: (1) Varying the payoffs, matching protocol, and pre-play communication, total surplus is consistently higher under centralization than decentralization. (2) Total surplus is maximized by combining centralization with free-form chat between a central manager and her subordinates. Total surplus is significantly higher in this treatment than for either a combination of decentralization and chat or centralization without chat. Both chat *and* centralization are necessary to

[1] See Lange and Taylor (1938) and Hayek (1945) for classic papers on centralization in an economy. See Mookherjee (2006) for a good survey of the older theoretical literature on centralization versus decentralization within a firm. Prominent examples of the more recent theory literature include Hart and Moore (2008), Alonso, Dessein, and Matouschek (2008a, 2008b and 2015), Rantakari (2008), Hart and Holmstrom (2010), Dessein, Garicano, and Gertner (2010). For recent empirical studies using observational data see Thomas (2011) and McElheran (2014).

achieve the best performance. (3) When chat and centralization are combined, lying is almost non-existent and the frequency of truth-telling does not respond to incentives. These results strongly contrast with typical findings from the experimental literature on lying. The resulting unambiguous transmission of information from divisions to the central manager explains why total surplus is maximized by the combination of centralization and chat.

Going into the details, our experimental environment is framed as a game between central management and two divisions of a firm. Either division managers choose their own product types (decentralization) or the central manager assigns product types to each division (centralization). Each division has a profit function that depends on local tastes, a randomly determined state of the world, and the chosen or assigned product types for the two divisions. The central manager is benevolent, earning the sum of the divisions' profits. The divisions know the current state of world, while central management only knows the probability distribution over states. Under centralization, the divisions make reports about the state of the world to the central manager prior to the central manager assigning product types. Unless otherwise noted, subjects play a series of games in fixed trios (i.e. "partners" matching).

Under decentralization, there exist multiple equilibria where the divisions coordinate by choosing a common product type. Differing local tastes give the divisions diametrically opposed preferences over these equilibria. Tastes depend on a randomly determined state of the world. The state of the world does not affect the divisions' ordinal preferences over equilibria, but instead determines which division cares most strongly about getting its preferred equilibrium. Efficiency, in the sense of maximizing total profits across the two divisions, requires coordinating on different product types over time, coming closer to the local tastes of one division or the other depending on the state of the world. *Ex ante*, the efficient equilibrium equalizes expected payoffs for the two divisions, but, *ex post*, the efficient equilibrium typically yields a higher profit to one division than the other. There also exists a "safe" equilibrium under decentralization that is safe (both divisions play their maximin strategy), does not require divisions to make different decisions as the state of the world changes, and yields both divisions identical payoffs in all states of the world. The safe equilibrium provides divisions with a relatively easy way to achieve coordination, unlike a battle-of-the-sexes game, but does not maximize total profits because it does not take advantage of the divisions' information about the state of the world.

Coordinating on the efficient equilibrium under decentralization is predicted to be challenging. There exists a large literature demonstrating that spontaneous coordination is unlikely to occur in the class of asymmetric coordination games that includes the decentralization game, the battle-of-the-sexes game (Cooper, DeJong, Forsythe, and Ross, 1989), and allocation games (Kuzmics, Palfrey and Rogers, 2014). The decentralization game differs from these other games because there exists an attractive, symmetric equilibrium. Even if the divisions manage to coordinate under decentralization, it is not obvious that the efficient equilibrium will emerge rather than the safe equilibrium.

Under centralization, the central manager's preferences are aligned with maximizing total profits across the two divisions. For any beliefs she might have about the state of the world, the central manager maximizes her profits by assigning the two divisions identical product types. The central manager therefore always chooses product types that match some equilibrium of the game under decentralization. The central manager would always choose the efficient outcome *if she knew the state of the world,* but the divisions have strong incentives to not reveal their information in their reports. The game is constructed so the only equilibria under centralization are babbling equilibria in which the division reports reveal no useful information to the central manager and the central manager makes product type choices equivalent to the safe equilibrium. Centralization is predicted to achieve coordination but not efficiency.

The preceding paragraphs illustrate the central tension in our experiments. Maximizing total surplus requires both coordinating divisions' product choices *and* taking advantage of their information. In theory decentralization can yield efficient coordination, but we anticipate that coordination, let alone efficient coordination, will prove difficult in practice. Centralization makes coordination trivial, but in theory efficiency cannot be achieved because the divisions' information will not be used.

Communication is pervasive in organizational settings and has been shown to improve efficiency in many environments (see discussion in Section 2). Anticipating that efficient coordination will not emerge in the absence of communication under either decentralization or centralization, we study behavior under a variety of communication structures. Communication, other than reports by divisions about the state of the world, is precluded in the two Baseline treatments (Baseline-Decentralization and Baseline-Centralization). We then study four treatments that add additional communication. Two combine structured communication, where subjects are limited to a small number of pre-specified messages, and decentralization: (1) Division Suggestions: Divisions send a series of messages proposing product choices before decisions are made and (2) Manager Suggestions: Prior to divisions choosing product types, the central manager sends a menu suggesting product types for each possible state of the world. The final two communication treatments feature free-form chat. In Chat-Decentralization, the two divisions can talk freely prior to making decisions. In Chat-Centralization, the central manager and two division managers can talk freely before the central manager chooses product types for the divisions. The theoretical predictions are unchanged for the structured communication and chat treatments, but related experimental results lead us to predict that pre-play communication, especially free-form chat, increases total surplus.

We focus on results from the second half of the experiment when play has settled down. In the Baseline treatments, centralization leads to higher total surplus than decentralization. Divisions struggle to coordinate their product types under decentralization. When the divisions do coordinate, it is primarily at the safe equilibrium which does not use their information. Centralization also fails to take advantage of the divisions' private information, albeit not quite in the manner predicted by the theory, but solves the coordination problem between divisions. Two follow-up treatments examine whether the results of the Baseline treatment are robust. The Strangers treatment checks whether the results depend on use of a "partners" design. Subjects keep the same role (division or central manager) throughout the experiment but are randomly matched into new trios in each round. The High State Losses treatment changes the payoff function to make the efficient equilibrium more attractive relative to the safe equilibrium. This gives decentralization a better chance of performing well relative to centralization. We find that neither the Strangers treatment nor the High State Losses treatment has much effect on outcomes relative to the Baseline treatment. Centralization yields higher total surplus than decentralization in both cases.

Neither treatment with a limited message space significantly increases total surplus over the Baseline-Decentralization treatment. In the Division Suggestions treatment, the problem of coordinating actions directly is replaced by the problem of reaching an agreement on what actions to take. Divisions often don't reach an agreement and, when they do, it is mostly on the safe equilibrium rather than the efficient equilibrium. In the Manager Suggestions treatment, suggesting play of the efficient equilibrium is quite effective. Unfortunately, central managers only suggest this about a third of the time. Even with communication structures designed to ease coordination, decentralization leads to frequent and persistent coordination failures.

Performance in the Chat-Decentralization treatment is very similar to performance in the Baseline treatment under *centralization*. Free-form chat eliminates the coordination problem under decentralization, but does so largely because divisions resort to the safe equilibrium. Total surplus is significantly higher in the Chat-Centralization treatment than either the Chat-Decentralization treatment or the Baseline-Centralization treatment. The Chat-Centralization treatment is the only case in which total surplus is significantly higher than the babbling equilibrium. Free-form chat increases total surplus under both decentralization and centralization, but it is not a magic bullet. The choice of organizational form matters even with free-form chat.

The combination of centralization and free-form chat works so well because it increases choice of the efficient outcome, making better use of the divisions' information. Central managers do *not* face the predicted problem of being unable to determine the true state of the world due to conflicting reports from the divisions. Instead, divisions almost never lie about the state of the world and there is rarely a conflict between the divisions' claims about the state of the world. Divisions are surprisingly non-strategic. They are no less likely to be truthful about the state of the world when lying would be advantageous and do not make use of partial lies. This contrasts strongly with the Baseline-Centralization treatment, where divisions' reports look like what would be expected from the extensive literature on lying (see discussion in Section 2): both lying and partial lying are common, and the probability of telling the truth versus lying responds strongly to incentives. "Fact-checking" is common in the Chat-Centralization treatment, with one division immediately correcting the other when it lies. We conjecture that divisions are reluctant to lie because of the potential embarrassment of being immediately contradicted by the other division. Central managers in the Chat-Centralization treatment generally get good quality information even in the difficult cases where the efficient equilibrium favors one division over the other and use this information to implement efficient outcomes.

The remainder of this paper is organized as follows. Section 2 summarizes related experimental papers and discusses how our work differs from the existing literature. Section 3 introduces the decentralization game and gives equilibrium predictions for the game. Section 4 lays out the experimental design and our *ex ante* hypotheses. The first part of Section 5 gives an overview of the treatment effects while the second part of the section digs into why these effects occur. Section 6 discusses the results and describes some directions for future work.

2. Related Literature: Evdokimov and Garfagnini (2018) and Hamman and Martínez-Carrasco (2018) directly compare centralization and decentralization. The first of these papers provides a direct experimental test of the theoretical models of Alonso et al. (2008a) and Rantakari (2008) in an experimental environment with two divisions and one central manager. The focus is on comparing the quality of horizontal communication between divisions with that of vertical communication between the divisions and central management, varying the importance of coordination. The results are largely in line with theoretical predictions: the quality of horizontal communication is significantly higher than that of vertical communication and divisions' actions are more coordinated when the importance of coordination is high. Even though the underlying games are related, the focus of our paper is different from Evdokimov and Garfagnini; they largely focus on testing predictions of Alonso et al. and Rantakari, while we are concerned with finding what type of communication is needed to achieve efficient coordination.

Hamman and Martínez-Carrasco (2018) examine decentralization in the context of a more complex environment involving task selection. Central management first chooses what type of workers to hire, homogenous or heterogeneous, and then chooses either to keep decision-making rights over task assignment (centralization) or to delegate task assignment to workers (decentralization). Decentralization is more common as task uncertainty grows, increasing the informational advantage of workers. There is a bias toward choosing centralization which they attribute to a desire for control by central managers. While the environment studied by Hamman and Martínez-Carrasco is rather different from ours, the endogenous choice of structures is an important innovation that we hope to explore in our environment in the future.

There are rich literatures in experimental economics dealing in isolation with coordination, truth-telling, and information processing. For general discussions of experiments related to coordination problems see Camerer (2003) and Cooper and Weber (forthcoming), and for surveys of experimental work on organizational issues see Camerer and Weber (2012) and Kriss and Weber (2013).

Most of the experimental literature focused on the intersection of coordination games and organizational economics uses weak-link games. Unlike our decentralization game, players in a weak-link game have aligned interests and asymmetric information does not play an important role.

The paper from the experimental literature on coordination most closely related to our current work is Cooper et al. (1989). They study the effect of communication in battle-of-the-sexes games. Without communication, coordination is difficult in these games due to the lack of a focal equilibrium. With one-way communication, coordination rates are high as the sender can call for her preferred equilibrium and the receiver generally follows. Two-way communication is less effective, although coordination rates improve somewhat with multiple rounds of pre-play communication.[2] If a pair sends messages that agree on an equilibrium, this equilibrium is generally played, but there are numerous cases where each calls for their preferred equilibrium, and play akin to the mixed-strategy equilibrium results.

Our Division Suggestions treatment is similar to the Cooper et al. treatment with multiple rounds of two-way pre-play communication, but there are two important differences. The safe equilibrium equalizes payoffs between the two divisions in all states of the world. Unlike in the battle-of-the-sexes game, this gives subjects in the decentralization game a simple way of coordinating that does not benefit one division over the other. Additionally, the Division Suggestions treatment is played using partners matching while Cooper et al. use strangers matching. The number of opportunities to communicate per round is the same in both experiments, but our experiment provides many more opportunities overall for a pair of divisions to communicate and reach an agreement. These two differences made us hope for a larger impact from two-way communication than the modest improvement observed by Cooper et al.[3]

There is also a large literature on subjects' willingness to lie (e.g. Gneezy, 2005; Erat and Gneezy, 2012; Fischbacher and Föllmi-Heusi, 2013; Abeler, Nosenzo, and Raymond, forthcoming; Gneezy, Kajackaite, and Sobel, 2018). Several striking regularities have emerged from this literature: (1) Subjects

[2] With no communication, the coordination rate is 48%. This improves to 95% with one-way communication as opposed to 55% with two-way communication. Allowing three rounds of two-way communication increases the coordination rate to 63%.

[3] We did not implement one-way communication, primarily because having only one division communicate is not particularly natural in the context of a decentralized organization. We doubt that one-way communication would increase coordination relative to two-way communication as much as in Cooper et al. In Cooper et al., the combination of strangers matching and role switching provides an obvious focal point in which players take turns getting their preferred equilibrium. Our experiment uses partners matching and roles are fixed. The resulting high degree of asymmetry makes it more difficult for subjects to coordinate on turn taking.

often tell the truth even when lying would pay more,[4] (2) the likelihood of lying is sensitive to incentives, and (3) partial lying (failing to either tell the truth or lie to the full extent that maximizes payoffs) is common. We observe these regularities in our Baseline, Strangers, and High State Losses treatments under centralization, but lying in the Chat-Centralization treatment is not sensitive to incentives and partial lies are rare. Theory has made a great deal of progress in understanding why individuals (at least partially) tell the truth and keep promises,[5] but we are not aware of any theory that explains the difference between the Chat-Centralization treatment and the other treatments with centralization as well as the existing literature.

There is a large literature showing that pre-play communication leads to great efficiency in social dilemmas (e.g. Dawes, MacTavish, and Shaklee, 1977; Isaac and Walker, 1988; Ostrom, Walker and Gardner, 1992; Charness and Dufwenberg, 2006; Cason and Mui, 2007; Cooper and Kühn, 2014) and coordination games with Pareto-ranked equilibria (e.g. Blume and Ortmann, 2007; Brandts and Cooper, 2007; Kriss, Blume, and Weber, 2016; Blume, Kriss, and Weber, 2017). In this literature, communication improves efficiency either by acting as a coordinating device or by enabling credible promises about intended actions. We observe that free-form chat increases efficiency under both decentralization and centralization. Under decentralization this stems from improved coordination, paralleling the existing literature, but under centralization improvement is driven by improved information transmission due to a near absence of lying by divisions. Our work identifies a new channel by which pre-play communication enhances efficiency.

There is ample evidence that the pro-efficiency effects of communication are greater with free-form chat than structured communication (e.g. Lundquist, Ellingsen, Gribbe, and Johannesson, 2009; Ben-Ner and Putterman, 2009; Charness and Dufwenberg, 2010; Cooper and Kühn, 2014; Brandts, Charness, and Ellman, 2016). We also observe that free-form chat is more effective than structured communication, although the games being studied and the mechanisms underlying the results differ from most of these studies. The closest to our work is Lundquist et al. who study a bargaining game with asymmetric information. They find that free-form messages reduce lying relative to pre-specified messages, although the effect is less extreme than what we observe and their study does not speak to partial lying. While the games are obviously different, comparing the results of Lundquist et al. with ours suggests that the near total absence of lies in the Chat-Centralization treatment is not solely due to the use of free-form chat since they still frequently observe deceptive lies (40%) with free-form messages.

---

[4] Studies of cheap talk games find a similar bias towards telling the truth (Cai and Wang, 2006; Sanchez-Pages and Vorsatz, 2007), which could stem from either an aversion to lying or a failure to grasp the strategic benefits of lying.

[5] Beyond the papers cited above, also see Charness and Dufwenberg (2006) on guilt aversion and Dufwenberg and Dufwenberg (2018) for a theory of partial lying.

Vespa and Wilson (2016) study information transmission in variations of the multi-sender cheap-talk model proposed by Battaglini (2002). Their experimental environment differs from ours in a critical aspect since, in keeping with Battaglini, they study games where messages should fully reveal the state of the world in equilibrium. We wanted to make the informational problems under centralization as severe as possible, so we designed a game where messages are *not* fully revealing (or even informative) in equilibrium. Nonetheless, their results have an obvious relationship to our experiments. They study three different games. In the two where it is relatively difficult (but possible) to infer the state of the world from senders' messages, receivers perform poorly at extracting information. The failure to extract information is even more severe in our Baseline, Strangers, and High State Losses treatments with centralization, as central managers make systematic errors in cases where the information extraction problem is rather simple.

3. The Decentralization Game: The "decentralization game" used in our experiments is not drawn from any specific model in the literature, but instead is intended to confront subjects with a challenging environment that accentuates the tradeoffs between centralization and decentralization. A simplified example gives the flavor of the game. Imagine a restaurant chain with two locations, one in Hamburger City and one in Pizza City. People in Hamburger City prefer hamburgers to pizza while those in Pizza City prefer pizza to hamburgers. The restaurants' earnings are determined by several factors. First, it is cheaper to buy ingredients in bulk, so costs are lower if both locations sell the same thing, either hamburgers or pizza. This leads to losses from coordination failure – each restaurant has higher costs if it sells something different from the other. Second, people in Hamburger City always want hamburgers, so the restaurant in Hamburger City always wants to sell hamburgers. Analogously, the restaurant in Pizza City always wants to sell pizza. This causes adaptation losses – losses if a restaurant sells products that do not match local tastes. Finally, tastes are subject to shocks. These don't affect the preferred option in either city, but instead affect the intensity of tastes. When hamburgers are fashionable, customers in Hamburger City strongly favor hamburgers. Customers in Pizza City still prefer pizza over hamburgers, but now don't care so much. This makes adaptation losses state dependent. A shock in favor of hamburgers causes the restaurant in Hamburger City to lose a lot if it sells pizza, while the restaurant in Pizza City loses little if it sells hamburgers. It follows that total payoffs across the two locations are maximized if both restaurants sell hamburgers. Similar logic applies when there is a shock in favor of pizza.

The two restaurants (divisions) are assumed to know the current state of the world (taste shock). The owner of the restaurant chain (central manager) only knows the distribution of taste shocks. Under decentralization, both divisions have the information needed to coordinate on a menu that maximizes total payoffs, but the restaurant in Hamburger City *always* wants to coordinate on selling hamburgers and the restaurant in Pizza City *always* wants to coordinate on selling pizza. There is no asymmetric information, but the divisions face a difficult coordination problem. With centralization the central manager can force the restaurants to coordinate but doesn't know whether hamburgers or pizza is the best choice to maximize total payoffs. She can ask the divisions, but they have an incentive to lie. Under centralization, coordination is trivial and asymmetric information becomes the central problem. This illustrates the main tension in the decentralization game: maximizing total surplus involves solving the coordination problem *and* taking advantage of the divisions' private information. Centralization is well-suited to achieving coordination while decentralization eases the informational problems. Which organizational form does better overall?

Table 1: Stage Game Payoffs (k1 = 54, k2 = 7, k3 = 4, and k4 = 14)

Game 1

	C1	C2	C3	C4	C5
R1	26, 54, 80	12, 29, 41	-2, 4, 2	-16, -21,-37	-30,-46,-76
R2	15, 40, 55	29, 43, 72	15, 18, 33	1, -7, -6	-13, -32,-45
R3	4, 26, 30	18, 29, 47	32,32, 64	18, 7, 25	4, -18, -14
R4	-7, 12, 5	7, 15, 22	21, 18, 39	35, 21, 56	21, -4, 17
R5	-18, -2, -20	-4, 1,-3	10, 4, 14	24, 7, 31	38, 10, 48

Game 3

	C1	C2	C3	C4	C5
R1	18, 46, 64	4, 29, 33	-10,12, 2	-24, -13, -37	-38, -38, -76
R2	15, 32, 47	29, 43, 72	15, 26, 41	1, 1, 2	-13, -24,-37
R3	12, 18, 30	26, 29, 55	40, 40, 80	26, 15, 41	12, -10,2
R4	1, 4, 5	15, 15, 30	29,26, 55	43,29, 72	29, 4, 33
R5	-10, -10, -20	4,1, 5	18,12, 30	32, 15, 47	46, 18, 64

Game 5

	C1	C2	C3	C4	C5
R1	10, 38, 48	-4, 21, 17	-18, 4,-14	-32,-13,-45	-46,-30, -76
R2	7, 24, 31	21, 35, 56	7, 18, 25	-7, 1,-6	-21,-16,-37
R3	4, 10, 14	18, 21, 39	32, 32, 64	18, 15, 33	4, -2, 2
R4	1,-4,-3	15, 7, 22	29, 18, 47	43, 29, 72	29, 12, 41
R5	-2,-18,-20	12,-7, 5	26,4, 30	40, 15, 55	54, 26, 80

Note: Each cell contains the payoffs for D1 (π_{D1}), D2 (π_{D2}), and CM (π_{CM}).

3.1. Stage Game Payoff Functions: There are three players in the game, a central manager (CM) and two division managers (D1 and D2). G denotes the state of the world: $G \in \{1,2,3,4,5\}$. The state of the world is randomly determined before players take any actions. Draws of G are i.i.d. with each state equally likely. To ease comparisons across treatments, the same draw of games was used in all sessions (although different groups in a session faced different draws). Both divisions are assumed to know the state of the world, but CM only knows the *ex ante* distribution over states. This represents a situation where both divisions know the general business conditions in the field, while central management does not. As standard nomenclature, we refer to the states of the world by the games they induce (e.g., Game 1 for G = 1). The divisions choose (under decentralization) or are assigned (under centralization) a product type from the space $T_i \in \{1,2,3,4,5\}$. Equations 1a, 1b, and 1c give the payoff functions for D1, D2, and CM respectively.

$$\pi_{D1} = k_1 - k_2|T_1 - 5| - k_3|T_1 - G| - k_4|T_1 - T_2| \quad \textbf{(Eq. 1a)}$$

$$\pi_{D2} = k_1 - k_2|T_2 - 1| - k_3|T_2 - G| - k_4|T_1 - T_2| \quad \textbf{(Eq. 1b)}$$

$$\pi_{CM} = \pi_{D1} + \pi_{D2} \quad \textbf{(Eq. 1c)}$$

For all treatments, $k_1 = 54$, $k_2 = 7$, and $k_4 = 14$. We set $k_3 = 4$ for all treatments except the High State Losses treatment. Table 1 displays the payoff tables for G = 1, 3, and 5, and copies of all five payoff tables with $k_3 = 4$ can be found in Appendix C. The three numbers in each cell of Table 1 correspond to the payoffs, denominated in ECUs, of D1 (π_{D1}), D2 (π_{D2}), and CM (π_{CM}). The row and column are the product types chosen by D1 and D2 respectively (or chosen for them by CM). The row (R) and column (C) numbers correspond to the product types chosen by the divisions (e.g. R3 $\equiv T_1 = 3$; C4 $\equiv T_2 = 4$).

The division payoffs are better understood through Equation 2. Division payoffs equal a constant minus three types of losses. The first two terms capture "taste losses," defined as losses due to deviating from local consumers' most preferred option. "Adaptation losses," given by the first term in the payoff function, are taste losses due to deviations from local tastes that do ***not*** depend on the state of the world. Divisions face diametrically opposed local tastes, with $T_1 = 5$ most preferred by D1's customers and $T_2 = 1$ most preferred by D2's customers. The second term in the division payoff functions represents "state losses," the component of taste losses which ***is*** state dependent. Taste shocks, represented by G, shift the degree to which the two divisions care about staying close to the product type that is most preferred at their location. If changing product types shifts a division away from its consumers' most preferred choice but towards the state of the world (current tastes), the marginal taste loss (sum of marginal adaptation loss and marginal state loss) is relatively low ($k_2 - k_3$). If the division shifts away from both its consumers' most preferred choice and the state of the world, marginal taste losses are high ($k_2 + k_3$). This implies that there is a kink in taste losses at the state of the world (G). The final term in the division payoffs represents "coordination losses," losses from not choosing the same product type as the other division. We assume $k_4 > k_2 > k_3 > 0$. These assumptions make the games induced by the five states of the world into coordination games where the two divisions have diametrically opposed interests.[6]

[6] If adaptation losses are higher than the sum of coordination and state losses ($k_2 > k_3 + k_4$), it is a dominant strategy for each division to perfectly conform to local tastes. For less extreme cases ($k_3 + k_4 > k_2 > k_4$), there are multiple equilibria but some choices become dominated implying that the subgames no longer have five equilibria and in some cases have a unique equilibrium. If state losses are larger than adaptation losses ($k_3 > k_2$), the divisions' preferences over equilibria are aligned and the coordination problem is largely eliminated.

$$\pi_D = k_1 - k_2 * adaptation\ loss - k_3 * state\ loss - k_4 * coordination\ loss \quad \textbf{(Eq. 2)}$$

As an example of how the payoffs work, suppose D1 chooses R1 and D2 chooses C1 for Game 1. This is the efficient (i.e., total surplus maximizing) equilibrium for Game 1, but is also the least preferred equilibrium for D1 yielding a payoff of 26 ECUs. If D1 deviates to choosing R2, D1 suffers 14 ECUs of coordination losses by ceasing to be coordinated with the other division and loses another 4 ECUs of state losses by moving away from the state of the world, but gains 7 ECUs in reduced adaptation losses by moving closer to his preferred choice R5. The total loss is 11 ECUs.

CM's payoff is the sum of the divisions' payoffs. In other words, central management seeks to maximize total firm profits. This need not be interpreted as benevolence on the part of central management, but instead can represent a setting where the rewards of division managers largely depend on how their division does while central management is concerned with profits across the entire firm.

Having five possible product types and five states of the world (rather than three) makes the game more complex, but nevertheless has distinct advantages. An important issue in our analysis is distinguishing whether play is consistent with an efficient (i.e., total payoff maximizing) or safe equilibrium. With only three states and a uniform distribution over states, the efficient and safe equilibria coincide for a third of the data. Using five states reduces the probability of being unable to distinguish the efficient and safe equilibria to one-fifth. Under centralization, having five product types makes it easier to detect partial lies.

The payoff functions in Equations 1a, 1b, and 1c are related to those used in Alonso et al.'s (2008a) model, but differ in several important ways. First, tastes in the Alonso et al. model are determined by two independent shocks (one for each division). Tastes in our model are subject to a single common shock. This accentuates the contrast between decentralization and centralization. Because there is no asymmetric information between D1 and D2, asymmetric information plays a role under centralization but not under decentralization. Second, the functional forms in Equations 1a and 1b are absolute values of differences rather than squared differences as used by Alonso et al. Using absolute values emphasizes the importance of coordination. Rather than shading their choices away from local tastes and towards the other division's choice, in equilibrium divisions do exactly the same thing as the other division. Finally, we assume each division manager is solely concerned with their own division's profits rather than a weighted average over both divisions' profits. Implicitly, this eliminates pay schemes where part of a division manager's pay is based on profit sharing across the entire firm. This design choice sharpens the incentives faced by division managers by making both coordination under decentralization and truthful revelation of divisions' private information under centralization more difficult. Our payoff functions reflect our primary goal in designing the decentralization game: confront subjects with a challenging environment that accentuates the tradeoffs between centralization and decentralization.

Following many recent papers on centralization (e.g. Alonso, Dessein, and Matouschek, 2008a, 2008b and 2015; Rantakari, 2008), we model the interaction between the central manager and divisions under centralization as a cheap talk game rather than a problem of mechanism design. Our various treatments change what types of communication are available, but do not fundamentally change the nature of the game from being a cheap talk game. Implicitly, we assume that the central manager cannot commit to a mechanism for eliciting information.

3.2: Equilibrium, Decentralization: Under decentralization, each division chooses a product type and CM is a passive bystander. Ignoring the payoff for CM, all five games are coordination games with five pure-strategy Nash equilibria where the two divisions choose the same product type: ($T_1 = T_2 = 1$), ($T_1 = T_2 = 2$), ($T_1 = T_2 = 3$), ($T_1 = T_2 = 4$), and ($T_1 = T_2 = 5$). For convenience, we refer to these outcomes as Equilibrium 1, Equilibrium 2, etc.

In all five games (see Table 1 and Table C1 in the appendix), there is a tension similar to the battle-of-the-sexes game since D1 most prefers Equilibrium 5 and least prefers Equilibrium 1, while the reverse is true for D2. Moreover, CM prefers the equilibrium that maximizes total surplus. This implies that the CM always wants a different equilibrium than at least one of the divisions and wants a different equilibrium than either division in Games 2 – 4. Alternative principles for selecting an equilibrium, such as safety and efficiency, suggest different ways of resolving the tension stemming from divisions' differing interests.

Unlike a battle-of-the-sexes game, the decentralization game offers a relatively easy way to coordinate and achieve equal payoffs since in all five games Equilibrium 3 yields the same payoff to both divisions. Equilibrium 3 is also safe, in the sense that $T_i = 3$ is the maximin strategy for both divisions in all five games. The problem with always playing the safe equilibrium is that, except in Game 3, Equilibrium 3 does not maximize total surplus.

All five games have an equilibrium that maximizes total surplus. This is always equivalent to the game number (i.e., Equilibrium 1 in Game 1, Equilibrium 2 in Game 2, etc.). The "efficient" equilibrium, where the divisions play the surplus-maximizing equilibrium in all states of the world, is procedurally fair (i.e., equalizes expected payoffs under the veil of ignorance about the state of the world; see Brandts et al., 2005) but yields asymmetric payoffs for all games except Game 3 once the state of the world is known. The efficient equilibrium is also relatively complex because the divisions must change their actions as the state of the world changes.

3.3: Equilibrium, Centralization: Under centralization, the two divisions do not choose rows and columns directly. After being informed about the state of the world (i.e., Game 1, Game 2, etc.) each division independently sends a message to the CM indicating which state of the world has been selected ($M_i \in \{1,2,3,4,5\}$). There is no requirement that these messages be truthful, a point emphasized in the instructions.[7] After receipt of the two messages, the CM chooses both a row and a column. The CM has no knowledge about which game has been selected beyond the initial distribution over states of the world and whatever information she gleans from the divisions' messages.

Conditional on enforcing coordination, Equations 1a and 1b are constructed so the sum of adaptation losses across D1 and D2 is a constant. This implies that the CM does not care about local tastes, but the divisions do. Given their opposing interests, the divisions have no incentive to be truthful with the CM. If both divisions always report the game where the efficient outcome is best for them (Game 5 for D1, Game 1 for D2), the best the CM can do is to choose the safe outcome ($T_1 = T_2 = 3$).[8] Any benefits from the divisions' private information are lost and the CM generally will not choose the efficient outcome.

More formally, we can prove the following theorem. Given that the CM must choose the same row and column ($T_1 = T_2$) in any Perfect Bayesian equilibrium (PBE), we henceforth refer to the CM as choosing a single product type in response to the divisions' messages.

Theorem: There does not exist a pure-strategy PBE where the CM chooses different product types for two different states of the world. This implies that the only pure-strategy PBE are babbling equilibria where the safe outcome ($T_1 = T_2 = 3$) is always chosen.

The lack of an informative equilibrium shown by this theorem does *not* reflect a generic property of cheap talk games with multiple senders. Battaglini (2002) proves that such games generically have an informative equilibrium when the state space is multidimensional. The decentralization game is constructed to give the two divisions diametrically opposed interests over a unidimensional state space. The resulting lack of an informative equilibrium makes information transmission impossible under centralization (in theory). This is in keeping with our goal of confronting subjects with a challenging environment that accentuates differences between centralization and decentralization.

[7] The instructions state " … [D1 and D2] will separately send messages to [CM] saying which game has been selected. This message can be truthful or not."

[8] Given that payoffs are linear, this isn't transparent. Define a CM's error as the difference between the product type she chooses (assume $T_1 = T_2$) and the efficient choice. Choosing the safe equilibrium limits the size of CM errors. If she chooses the safe equilibrium, her average error is 1.2. If she chooses product type 2 or 4, the average error rises to 1.4. Choosing 1 or 5, the average error goes up to 2.0.

The theory assumes messages are cheap talk, with the divisions incurring no cost, pecuniary or psychological, for sending false messages. If we add a psychological cost for sending false messages (e.g. Kartik, 2009), it is trivial to construct cases where truthful revelation is consistent with an equilibrium. For example, let $c_L*|M_i - G|$ be a division i's psychological cost of lying. If $c_L > k_2 - k_3$, there exists an equilibrium in which both divisions truthfully reveal their information.[9]

4. Experimental Design and Hypotheses: Subjects played 18 rounds in all treatments. They were assigned the role of CM, D1, or D2 at the beginning of the session and these roles remained constant throughout the session. In treatments with decentralized decision-making, the participants in the CM role were pure observers. We did this to keep the possible influence of other-regarding preferences constant across treatments. There were ten different treatments as described below. We used a between-subjects design, so each subject participated in just one of the treatments.

Baseline (Baseline-Decentralization and Baseline-Centralization): This was the baseline treatment where subjects played the game described in Section 3 without any additional communication. We set $k_3 = 4$, yielding the payoff tables shown in Table 1. This is a "Goldilocks" value for state losses, high enough that divisions care about the state of the world but not so high that adaptation losses are not clearly a larger concern than state losses. Partners matching was used (i.e. participants were matched with the same two other subjects throughout the entire experiment). At the beginning of the session, they were told that the matching would be fixed for the first nine-round block and were given no information about the second block. After the first block was over, subjects were informed that there would be an additional nine rounds of the same game using the same groups as in the first nine rounds.

Subjects received feedback about the realized state of the world (i.e. the game being played) and the product types chosen after each round. In the centralization treatment, this made it possible for CMs to know if a division had lied about the game being played.

We conducted three sessions with decentralization (i.e., divisions choose their product types) and three with centralization (i.e., divisions send messages to the central manager who then chooses their product types). There were nine three-person groups per session, giving a total of 27 subjects per session and 27 groups per treatment.

Strangers (Strangers-Decentralization and Strangers-Centralization): This treatment was identical to Baseline except for use of strangers matching; that is, groups changed from round to round. The matching

[9] If both divisions send the same message, the CM chooses the corresponding equilibrium. If $M_i = 1$ and $M_j = 2$, where $i, j \in \{1,2\}$ and $i \neq j$, the CM chooses Equilibrium 2. If $M_i = 4$ and $M_j = 5$, where $i, j \in \{1,2\}$ and $i \neq j$, the CM chooses Equilibrium 4. Otherwise, the CM chooses Equilibrium 3.

was constructed so no participant met another person twice in a nine-round block (a point which the instructions stressed). At no time were participants informed about the identities of the other two people in their group. Because groups were not independent within a session, we conducted five sessions per treatment rather than three. There were 27 subjects in each session.

Our design focused on partners matching as the natural case since we are interested in the effect of organizational structure within long-lasting organizations. The Strangers treatment is a robustness check, testing whether the results of the Baseline treatment were sensitive to what type of matching was used.

High State Losses (HSL-Decentralization and HSL-Centralization): This treatment was identical to the Baseline treatment except we increased state losses ($k_3 = 6$ vs. $k_3 = 4$). Increasing k_3 does not affect the theoretical predictions for the game under either decentralization or centralization, but minimizes the difference between adaptation and state losses. This reduces the tension between divisions since the gain for moving from the efficient equilibrium to a division's most preferred equilibrium is tiny. For example, moving from Equilibrium 1 to Equilibrium 5 in Game 1 gains D1 12 ECUs in the Baseline treatment, but only 4 ECUs in the High State Losses treatment.

The High State Losses treatment is a second robustness check. Previous experiments with asymmetric coordination games (e.g. the battle-of-the-sexes game) suggest that achieving coordination, let alone efficiency, will be challenging under decentralization. Even if coordination occurs, the tension between divisions makes the safe equilibrium attractive, sacrificing efficiency in order to achieve coordination. Under centralization, achieving efficiency is difficult because the tension between divisions provides a strong incentive to deceive the CM. The High State Losses treatment weakens tension between the divisions, letting us explore how the results change when the environment is less challenging.

The remaining four treatments explore the effects of adding additional cheap talk to the game. The first two, Division Suggestions and Manager Suggestions, involve structured communication. Subjects can only send a small number of pre-specified messages. The final two treatments, Chat-Decentralization and Chat-Centralization, use free-form chat, allowing subjects an unlimited message space.

Division Suggestions: This treatment was identical to the Baseline-Decentralization treatment, except pre-play communication between divisions was added. Prior to the divisions' choices of product types, each game began with three rounds of messages. Within each round of messages, the divisions simultaneously chose a pair of messages suggesting product types for themselves and the other division. Divisions observed each other's messages at the end of each round of messages. The purpose of having three rounds of messaging (rather than one) was to make it easier for divisions to agree upon a course of action.

Manager Suggestions: This treatment was identical to the Baseline-Decentralization treatment, except the CM sent a message to the divisions prior to each round of play. This message suggested product types for both divisions in each of the five possible games. The full message (a 5 x 2 matrix) was shown to both divisions prior to making choices. The divisions knew that both received identical messages.

Chat (Chat-Decentralization and Chat-Centralization): These two treatments parallel the corresponding Baseline treatments except for the nature of pre-play communication. Under decentralization, the divisions had two minutes to engage in free-form discussions via chat before choosing product types. They could discuss whatever they chose. In practice, discussions largely focused on the obvious topic, how to play the game. The CM saw the discussion but could not participate. Under centralization, the structured messages about the state of the world were eliminated and replaced by a free-form discussion between the divisions and the CM. While the divisions were not specifically instructed to share information about the state of the world, this was a natural and typical topic of conversation (making the structured messages redundant). All three players saw all messages and again there was a two-minute time limit.

The two treatments with structured communication (Division Suggestions and Manager Suggestions) focus on specific aspects of communication under decentralization. Both provide coordination mechanisms, but of different sorts. The Division Suggestions treatment gives the divisions an opportunity to coordinate choices prior to picking product types. This makes the agents responsible for choosing product types directly involved in achieving coordination, but doesn't resolve the problem that coordination may be difficult due to their divergent interests. The Manager Suggestions treatment allows the CM to act as a leader. She cannot impose outcomes, since decision making is decentralized, but can act as a neutral arbiter, using suggestions to create a focal point for the two divisions.

The two chat treatments eliminate the restrictions that naturally accompany use of a limited message space. With chat, players can make unlimited asynchronous proposals and proposals can extend across multiple rounds. Proposals can be accompanied by an explanation. The CM (under centralization) can appeal to the better nature of the divisions in an attempt to get them to be truthful.

Each session began with instructions (see Appendix B). Participants had printed copies of the payoff tables for all five games. The sessions were run at the LINEEX lab at the University of Valencia, with undergraduate students from the university as participants. The payoffs were denominated in Experimental Currency Units, with 1 ECU = 0.2€. Participants were paid their cumulative earnings for all rounds. Including a 5€ show-up fee, average pay was 19.90€. Each session lasted around an hour.

Ex ante, we had six hypotheses about the data. The first four concern the treatments that do not add pre-play communication of some type to the game (Baseline, Strangers, and High State Losses). The efficient equilibrium under decentralization achieves the maximum possible total surplus. No equilibrium that achieves the maximum possible total surplus exists under centralization. This leads to H1. Of course, there are *many* possible equilibria under decentralization and achieving coordination, let alone the efficient equilibrium, may be difficult. So there is reason *ex ante* to be skeptical about H1.

H1: *Total surplus will be greater under decentralization than under centralization in the Baseline, Strangers, High State Losses treatments.*

The theory implies that messages under centralization, consistent with play of a babbling equilibrium, will contain no useful information. H2 follows from this logic. Once again, there are good reasons to be skeptical. Our design differs from many of the existing experiments, especially since there is more than one sender, but the general finding that individuals are reluctant to tell lies is likely to apply.

H2: *Under centralization, messages from the divisions will contain no useful information about the state of the world. Total surplus will not exceed the expected payoff from a babbling equilibrium.*

The set of (subgame perfect) equilibrium outcomes is largely unaffected by the type of matching under either decentralization or centralization.[10] We nevertheless expected strangers matching to reduce performance in both cases. Turn-taking is natural for divisions in a fixed group due to reciprocity; if the divisions coordinate at a stage-game equilibrium that gives one division a higher payoff, that division can expect the favor to be returned in the future. Partners matching also makes it easier for a central manager to learn which divisions are truthful. This is not an issue in the theory since divisions are predicted to babble, but matters in practice. These advantages of having fixed groups lead to H3.

H3: *Under either decentralization or centralization, total surplus will be lower in the Strangers treatment than in the Baseline treatment.*

Increasing state losses does not affect the theoretical predictions, but changes the marginal costs of deviating from the efficient equilibrium as well as the cost of telling lies. These changes should favor efficiency under both decentralization and centralization, yielding H4.

H4: *Under either decentralization or centralization, total surplus will be higher in the High State Losses treatment than in the Baseline treatment.*

Turning to the treatments that add structured pre-play communication (Division Suggestions and Manager Suggestions), the theoretical predictions are not changed by either type of pre-play

[10] With partners matching it is possible to construct subgame perfect equilibria in which outcomes that are not stage-game equilibria are played in early rounds (if one division deviates, they are punished via play of the other division's ideal equilibrium). Such equilibria are Pareto-dominated and implausible given the complexity involved. Under centralization, the stage game has a single payoff vector consistent with equilibrium. Backward induction therefore implies that the stage-game equilibrium is played in every round with partners matching.

communication.[11] Cooper et al. (1989) only observed modest improvements from a treatment analogous to Division Suggestions, but there are differences in the structure of our experiment and game that gave us reason to be more optimistic (see Section 2 for a detailed discussion of these differences). Turning to the Manager Suggestions treatment, suggestions from the experimenter or an external leader can act as a coordination device in weak-link coordination games (e.g., Van Huyck, Gillette, and Battalio, 1992; Brandts and Cooper, 2007; Brandts, Cooper, and Weber, 2015). We conjectured that the CM's messages would have a similar effect in the decentralization game, but there are reasons to be skeptical since the decentralization game, unlike weak-link games, does not have an efficient symmetric equilibrium. The preceding conjectures are summarized in H5. This hypothesis is stated relative to the Baseline-Decentralization treatment, as both limited message space treatments modify this treatment, but we hoped that both treatments would also yield higher total surplus than the Baseline-Centralization treatment. We did not have any *ex ante* prediction comparing the two structured communication treatments.

H5: *We expect higher total surplus in the Division Suggestions and Manager Suggestions treatments than in the Baseline-Decentralization treatment.*

The final hypothesis covers the chat treatments. Cooper and Kühn (2014) find that free-form communication outperforms structured communication, largely by improving coordination on an efficient equilibrium. While the games are quite different, we expect that the ability to make unlimited asynchronous proposals along with the ability to explain proposals will similarly improve coordination under decentralization. We therefore expect Chat-Decentralization to lead to higher surplus than the Division Suggestions and Baseline-Decentralization treatments. For similar reasons we expect Chat-Centralization to yield higher surplus than Baseline-Centralization. We do not have any ex ante predictions with respect to the comparison between the two chat treatments. Our conjectures are summarized in H6.

H6: *Chat-Decentralization will lead to higher total surplus than Division Suggestions and Baseline-Decentralization. Chat-Centralization will lead to higher surplus than Baseline-Centralization.*

5. Results: Section 5.1 gives an overview of the main treatment effects, and Section 5.2 digs into the process underlying the treatment effects.

5.1. Treatment Effects: Unless otherwise noted, statistical tests comparing treatments are Wilcoxon rank-sum tests and comparisons with the babbling equilibrium are Wilcoxon matched-pairs signed-rank tests.

[11] Using the terminology of Aumann (1990), a message from a division in the Division Suggestions treatment calling for play of the efficient equilibrium is self-committing but not self-signaling. There exist equilibria in the Manager Suggestions treatment where the divisions' actions are conditioned on the CM's message, but this does not expand the set of possible equilibrium outcomes.

The unit of observation is a single group except for the Strangers treatments where it is a single session. ***Statistical comparisons are based on the second half of the experiment (Rounds 10 – 18) when play has had a chance to settle down.*** An observation is the average value of the variable in question over these rounds. Surpluses from the babbling equilibrium are based on what the total surplus would have been if the CM had always set $T_1 = T_2 = 3$ as predicted.

Table 2: Summary of Outcomes

Rounds 1 - 9

Treatment	All Data		Games 1, 2, 4, and 5		
	Total Surplus	% Coordinate	% Safe Equilibrium	% Efficient Equilibrium	% Other Equilibrium
Baseline-Decentralization	53.5	46.1%	27.0%	9.0%	5.8%
Baseline-Centralization	70.0	97.5%	28.6%	38.1%	31.7%
Strangers-Decentralization	50.0	38.3%	19.0%	12.7%	3.5%
Strangers-Centralization	69.3	97.8%	28.9%	31.7%	37.5%
HSL-Decentralization	49.2	43.6%	16.9%	18.0%	3.2%
HSL-Centralization	66.5	95.9%	24.3%	40.7%	30.7%
Division Suggestions	62.7	72.8%	44.4%	20.1%	5.3%
Manager Suggestions	57.9	55.1%	23.3%	23.8%	3.7%
Chat-Decentralization	64.2	82.3%	36.5%	26.5%	16.9%
Chat-Centralization	71.3	98.4%	40.2%	33.3%	24.3%

Rounds 10 - 18

Treatment	All Data		Games 1, 2, 4, and 5		
	Total Surplus	% Coordinate	% Safe Equilibrium	% Efficient Equilibrium	% Other Equilibrium
Baseline-Decentralization	61.4	69.5%	48.1%	12.6%	4.9%
Baseline-Centralization	71.7	99.6%	31.1%	36.1%	32.2%
Strangers-Decentralization	55.0	46.7%	31.5%	7.2%	3.0%
Strangers-Centralization	70.3	99.5%	37.7%	28.9%	33.4%
HSL-Decentralization	60.0	59.3%	11.5%	40.4%	1.1%
HSL-Centralization	68.2	97.5%	24.6%	48.1%	24.0%
Division Suggestions	65.7	77.8%	36.1%	30.6%	8.2%
Manager Suggestions	64.4	69.5%	32.8%	30.1%	4.4%
Chat-Decentralization	72.2	97.5%	41.0%	36.1%	20.8%
Chat-Centralization	75.2	100.0%	34.4%	50.8%	14.8%

Table 2 summarizes outcomes by treatment, with the two sub-tables dividing the data between the first (Rounds 1 – 9) and second (Rounds 10 – 18) halves of the experiment. The first two columns show average total surplus and coordination rates (i.e., choice/assignment of identical product types: $T_1 = T_2$) across all games. As a point of comparison, recall that the maximum possible total surplus always equals 80. The final three columns look at the frequency of specific outcomes in Games 1, 2, 4, and 5. As defined previously, the safe equilibrium refers to mutual choice of 3 ($T_1 = T_2 = 3$) and the efficient equilibrium indicates choices matching the state of the world ($T_1 = T_2 = G$). "Other" equilibrium refers to any other outcome where the divisions' product choices are identical. These three columns do not use data from Game 3 because the efficient and safe equilibria coincide in this case.

We begin by comparing total surplus under centralization and decentralization in the Baseline, Strangers, and High State Losses treatments.[12] Total surplus is always higher under centralization than

[12] We define total surplus as the sum of the payoffs for D1 and D2. This is equivalent to CM's payoff.

decentralization. The difference between centralization and decentralization narrows over time, but is significant over Rounds 10 – 18 in all three cases: Baseline (n = 54; z = 3.89; $p < .01$), Strangers (n = 10; z = 2.61; $p < .01$), and High State Losses (n = 54; z = 2.89; $p < .01$).

The low total surplus under decentralization relative to centralization stems from frequent failures to coordinate on identical product types. For all three centralization treatments, coordination is close to 100% throughout; subjects playing as CMs have no difficulty in understanding that choosing the same product types for both divisions is always preferable. In stark contrast, coordination in the decentralization treatments is below 50% in the first half of the experiment. This increases with experience but never gets close to the near-perfect coordination observed under centralization. The difference in coordination rates between decentralization and centralization is significant across Rounds 10 – 18 in all three settings: Baseline (n = 54; z = 4.82; $p < .01$), Strangers (n = 10; z = 2.66; $p < .01$), and High State Losses (n = 54; z = 5.91; $p < .01$).

Result 1: *Total surplus is higher under centralization than decentralization for all three settings (Baseline, Strangers, and High State Losses). The data are not consistent with H1. Lower coordination rates largely explain the lower total surplus under decentralization.*

Even when coordination occurs under decentralization, the divisions' product choices often fail to use their information. Coordination at the safe equilibrium represents 75% of the cases in Rounds 10 – 18 where the divisions coordinate and the efficient and safe equilibria do *not* coincide ($G \neq 3$) in Baseline-Decentralization and Strangers-Decentralization. The High State Losses treatment is designed to make coordination at the efficient equilibrium more attractive, and indeed the efficient equilibrium is played in 76% of the cases where the divisions coordinate (Rounds 10 – 18, $G \neq 3$) in HSL-Decentralization. This is only a limited success, as the divisions still fail to coordinate or coordinate at inefficient outcomes in the majority of cases (50% across all five games in Rounds 10 – 18).

Result 2: *When divisions coordinate under decentralization, play is generally consistent with the safe equilibrium in the Baseline and Strangers treatments, implying a failure to use the divisions' information. Play consistent with the efficient equilibrium is more frequent in the HSL-Decentralization treatment.*

The difference between total surplus under centralization and the babbling equilibrium is small in all three settings – recall that the expected total surplus from the babbling equilibrium is either 70.4 (Baseline and Strangers) or 65.6 (High State Losses). Differences from the babbling equilibrium are not significant for Baseline-Centralization (n = 27; z = 0.78; $p > .10$) or Strangers-Centralization (n = 5; z = 1.48; $p > .10$). In HSL-Centralization the total surplus is weakly greater than the babbling equilibrium (n = 27; z = 1.89; $p < .10$). Underlying this is greater play of the efficient equilibrium in HSL-Centralization than either Baseline-Centralization or Strangers-Centralization, although these differences are not statistically significant in either case.[13] Easing the tension between the two divisions' interests somewhat improves usage of the divisions' information.

Given the high degree of coordination under centralization and average total surpluses close to the expected value from the babbling equilibrium, it is tempting to conclude that play under centralization converges to the babbling equilibrium as predicted by the theory (implying that CMs never make use of the divisions' information). This is not the case. In the babbling equilibrium, the safe equilibrium (Equilibrium 3) should always be chosen. Returning to Table 2, the frequency of the safe equilibrium for cases where the safe and efficient equilibria can be distinguished ($G \neq 3$) rises slightly over time but remains below a third (32.3%) averaging across the three centralization treatments for Rounds 10 – 18.

Result 3: *Total surplus under centralization is about the same as the babbling equilibrium. The data are consistent with H2 along this dimension.*

The matching protocol has little effect (holding the payoff function fixed). Total surplus is lower in Strangers-Decentralization than Baseline-Decentralization, but this difference is not significant ($n = 32$; $z = 1.48$; $p > .10$). Coordination is less frequent in Strangers-Decentralization than Baseline-Decentralization, but this difference also misses statistical significance ($n = 32$; $z = 1.55$; $p > .10$). As should be obvious from Table 2, neither total surplus nor the frequency of coordination differs by matching protocol under centralization. We have previously shown that the difference between centralization and decentralization is significant for either partners or strangers matching. Overall, the partners-strangers distinction is not crucial in our data.

Result 4: *Total surplus is not significantly lower in the Strangers treatment than in the Baseline treatment under either decentralization or centralization. The data do not support H3.*

Under either centralization or decentralization, total surplus is lower in the High State Losses treatment than the Baseline treatment. The difference narrows over time and is not significant in Rounds 10 – 18 for either decentralization ($n = 54$; $z = 0.59$; $p > .10$) or centralization ($n = 54$; $z = 1.41$; $p > .10$). The data do not support our prediction of higher total surplus in the High State Losses treatment.

Coordination rates are lower in HSL-Decentralization than Baseline-Decentralization, but the difference is modest and not significant ($n = 54$, $z = 1.53$; $p > .10$). We predicted the High State Losses treatment would encourage play of the efficient equilibrium and, as noted above, it does so. So why is H4 not supported by the data? The modest increase in coordination failure in HSL-Decentralization more than offsets the movement toward the efficient equilibrium.

[13] Baseline-Centralization vs. HSL-Centralization: $n = 54$; $z = 1.04$; $p > 0.10$. Strangers-Centralization vs. HSL-Centralization: $n = 32$; $z = 1.23$; $p > 0.10$

Result 5: *Total surplus is lower, albeit not significantly, in the High State Losses treatments than the Baseline treatments. The data are not consistent with H4.*

Summarizing our findings thus far, centralization yields higher total surplus than decentralization but is no magic bullet, doing no better than predicted for the babbling equilibrium. CMs almost always enforce coordination on a single product type but often fail to choose the efficient outcome. This implies that CMs do not take advantage of the divisions' information. A similar issue arises under decentralization, as coordination is often at the safe equilibrium, with the added problem that coordination failure is common. This leads to the question addressed by the remaining four treatments: do added opportunities to communicate help subjects achieve coordination *and* take advantage of the divisions' information?

Returning to Table 2, we compare total surplus in the four communication treatments with the Baseline treatment both under decentralization and under centralization. The Baseline treatment is the correct point of comparison since it uses the same payoffs ($k_3 = 4$) and the same matching scheme (partners) as the four communication treatments.

Neither structured communication treatment (Division Suggestions and Manager Suggestions) has much impact. Both yield higher total surplus than Baseline-Decentralization, but the difference narrows in later rounds and neither difference is significant across Rounds 10 – 18 (Baseline-Decentralization vs Division Suggestions: $n = 54$; $z = 1.43$; $p > 0.10$; Baseline-Decentralization vs Manager Suggestions: $n = 54$; $z = 0.81$; $p > 0.10$). Neither treatment does as well as Baseline-Centralization. The gap narrows with experience, but both differences are significant across the second half of the experiment, albeit weakly for Division Suggestions (Baseline-Centralization vs Division Suggestions: $n = 54$; $z = 1.92$; $p < 0.10$. Baseline-Centralization vs Manager Suggestions: $n = 54$; $z = 2.61$; $p < 0.01$).

The Division Suggestions and Manager Suggestions treatments have little effect on total surplus because neither does much to improve the coordination rate relative to Baseline-Decentralization (Baseline-Decentralization vs Division Suggestions: $n = 54$; $z = 0.96$; $p > 0.10$; Baseline-Decentralization vs Manager Suggestions: $n = 54$; $z = 0.07$; $p > 0.10$). These treatments do generate some improvement in moving play away from the safe equilibrium and toward the efficient equilibrium. Excluding data from Game 3 where the efficient and safe equilibria cannot be distinguished, this improvement is weakly significant for Division Suggestions ($n = 54$; $z = 1.75$; $p < .10$) but not Manager Suggestions ($n = 54$; $z = 1.31$; $p > .10$).

Result 6: *The Division Suggestions and Manager Suggestions treatments do not yield higher total surplus than Baseline-Decentralization, and do significantly worse than Baseline-Centralization. The data do not support H5.*

Both treatments with free-form chat yield higher average surplus than Baseline-Decentralization (Baseline-Decentralization vs Chat-Decentralization: $n = 54$; $z = 4.14$; $p < .01$; Baseline-Decentralization vs Chat-Centralization: n = 54; z = 5.19; p < .01). Chat-Decentralization also beats the two structured communication treatments as hypothesized (Division Suggestions vs Chat-Decentralization: n = 54; z = 2.30; p < .05; Manager Suggestions vs Chat-Decentralization: n = 54; z = 2.87; p < .01).

Chat-Decentralization yields almost exactly the same average total surplus as Baseline-Centralization (n = 54; z = 0.44; p > .10). Like centralization, free-form chat succeeds under decentralization by largely eliminating coordination failure, but the likelihood of the efficient equilibrium is identical in Rounds 10 – 18 for Chat-Decentralization and Baseline-Centralization. This implies that usage of the divisions' information is poor in Chat-Decentralization. Consistent with this, total surplus is not significantly higher than predicted by the babbling equilibrium (n = 27; z = 1.61; p > .10).

Chat-Centralization yields significantly higher total surplus than Baseline-Centralization (n = 54; z = 2.94; p < 0.01) or Chat-Decentralization (n = 54; z = 2.45; p < 0.05). Not only is coordination 100% perfect across Rounds 10 – 18, but play of the efficient equilibrium is increased relative to either Baseline-Centralization (n = 54; z = 2.20; p < 0.05) or Chat-Decentralization (n = 54; z = 1.86; p < 0.10). This increase in play of the efficient equilibrium allows Chat-Centralization to yield significantly higher total surplus than predicted by the babbling equilibrium (n = 27; z = 3.65; p < .01). The combination of centralization and free-form chat yields higher total surplus by using the divisions' information better.

A good measure of how much the various treatments improve total surplus over the babbling equilibrium is the percentage of possible gains over the babbling equilibrium that are achieved. This is defined as the difference between average total surplus over Rounds 10 – 18 and the average total surplus predicted by the babbling equilibrium divided by the maximum possible gain over the babbling equilibrium.[14] Baseline-Centralization and Chat-Decentralization only achieve a small fraction of the possible gains over the babbling equilibrium (6.5% and 12.3% respectively), but Chat-Centralization gains half of the possible gains over the babbling equilibrium (50.4%).

Result 7: *Our results confirm H6 as Chat-Decentralization outperforms Division Suggestions and Manager Suggestions, and Chat-Centralization outperforms Baseline-Centralization. Chat-Centralization also yields higher total surplus than either Chat-Decentralization or the babbling equilibrium.*

To summarize, centralization yields higher total surplus than decentralization in the four "apples-to-apples" comparisons (Baseline, Strangers, High State Losses, and Chat). Structured communication has little effect on performance relative to the Baseline treatment, but free-form chat leads to higher total surplus

[14] The maximum possible total surplus is 80. Given the games drawn in Rounds 10 – 18, the expected payoff from the babbling equilibrium is 71.11. The maximum possible gain over the babbling equilibrium is therefore 8.89. The gains achieved relative to the babbling equilibrium are given by the difference between the average total surplus in Rounds 10 – 18 and 71.11. For instance, the gains achieved relative to the babbling equilibrium for Baseline-Centralization equal 0.58 (71.69 – 71.11). The percentage of possible gains achieved equals 0.58/8.89 = .065 or 6.5%.

under either decentralization or centralization. Chat-Centralization does especially well, yielding higher total surplus than the babbling equilibrium due to increased play of the efficient equilibrium.

5.2. Process: Section 5.2 examines the processes underlying the treatment effects described in Section 5.1. We focus on three specific questions: (1) Why is performance in the three centralization treatments without pre-play communication (Baseline-Centralization, Strangers-Centralization, and HSL-Centralization) no better than predicted by the babbling equilibrium? (2) Why do the two treatments with structured communication (Division Suggestions and Manager Suggestions) have little effect compared to Baseline-Decentralization? (3) Why is performance in Chat-Centralization better than either Baseline-Centralization, Chat-Decentralization, or the babbling equilibrium prediction?

5.2.a. Information Transmission: As described in the discussion leading up to Result 3, total surplus in the three centralization treatments without additional pre-play communication (Baseline-Centralization, Strangers-Centralization, and HSL-Centralization) is little better than predicted by the babbling equilibrium, but subjects do ***not*** play a babbling equilibrium. This subsection studies the details of what messages divisions send and how central managers interpret these messages, with the goal of understanding why centralization fails to produce higher surplus than the babbling equilibrium in these three treatments.

Two things have to happen under centralization for a group to take advantage of the divisions' information. The divisions have to send messages that are informative about the state of the world, and the central manager has to correctly interpret the information contained in the messages. The theory presented in Section 3 focuses on the first issue and concludes that information transmission will fail since the divisions have no incentive to send informative messages. Built into the theory is an assumption that the messages would be interpreted correctly if informative. In reality, messages sent by D1 and D2 contain useful information, but CMs make systematic errors in using messages. Total surplus is about the same as predicted by the babbling equilibrium because the advantages from better than expected information transmission are wiped out by errors in using this information.

If messages are uninformative, as the theory predicts, there should be no correlation between messages and the game being played. Instead, pooling data across all three treatments and all eighteen rounds, there is strong correlation between the message sent and the game being played ($\rho = .37$). Play of a babbling equilibrium implies that divisions will only tell the truth in 20% of the observations, but the observed likelihood of truth-telling is 55%. Even when it is least beneficial to do so (Game 1 for D1 or Game 5 for D2), 24% of messages tell the truth. Truth-telling decreases only slightly with experience, falling from 56% in Rounds 1 – 9 to 54% in Rounds 10 – 18. This implies that truth-telling is probably *not* due to a failure to grasp the strategic value of lying. Purely self-interested division managers should always send a message corresponding to their most preferred equilibrium. This is the most common type of self-serving lie in all three treatments, but 38% of self-serving lies are partial lies (i.e. the message lies strictly between the true game and the division manager's preferred equilibrium). This resembles the partial lying observed in papers like Fischbacher and Föllmi-Heusi (2013).[15]

Regressions reported in Table D1 in Appendix D provide statistical support for the preceding discussion. This analysis shows that the relationship between the game being played and the message sent is statistically significant, stable over time, and significantly stronger in the High State Losses treatment.

Result 8: *The messages sent by division managers are informative. The data are not consistent with the second part of H2.*

Table 3: CM Choices as a Function of Messages

		Message D2				
		1	2	3	4	5
Message D1	1	1.35	---	---	---	---
	2	2.00	2.23	---	---	---
	3	2.46	2.69	2.98	---	---
	4	2.79	3.14	3.55	3.75	---
	5	3.11	3.24	3.73	4.12	4.74

On aggregate, the CMs seem to realize that the division managers' messages contain useful information. Table 3 shows the CM's average choices as a function of the messages sent by the two division managers. Once again, data are pooled across treatments and all eighteen rounds. Cells where D2's message is greater than D1's message are left blank due to the small number of observations. When the two messages coincide the CM follows the messages closely. When the two messages differ, the CM's choices increase in each division's message (holding the other division's message fixed). Hence, CMs consistently respond to the divisions' messages. Table D2 in Appendix D reports regressions providing more formal analysis of the relationship between DMs' messages and the CM's choices. This analysis confirms that CMs respond significantly to the divisions' messages, with no significant difference between the three treatments.[16]

Although CMs respond to the information contained in divisions' messages, it doesn't follow that their use of this information is optimal. CMs often respond to the divisions' messages in ways that are hard to explain. For example, in the frequent cases (29% of all observations) where the two divisions' messages

[15] There are a small number of observations (3.7% of all messages) where divisions tell lies that are not self-serving (not shaded in the direction of the division's preferred equilibrium). The proportion of such lies falls with experience (4.9% in Rounds 1 – 9 vs. 2.6% in Rounds 10 – 18), suggesting that these are primarily errors.

[16] The regressions also examine CMs' reactions to truth-telling in the previous round. CMs are significantly more sensitive to a message from a division manager who told the truth in the previous round.

match ($M_1 = M_2$), they are almost certainly telling the truth (97%). Not surprisingly, it is a best response to assign both divisions the product type that corresponds to their messages ($T_1 = T_2 = M_1 = M_2$).[17] CMs do *not* play the best response in 19% of these observations. CMs making this type of error earn an average payoff of 62.7 ECUs, compared with 79.3 ECUs for those who play the best response. Failure to play the best response when the divisions' messages agree drops but does not vanish with experience, dropping from 23% in the first half of the experiment to 15% in the second half.

Another common type of error occurs when D1 and D2 send diametrically opposed messages by choosing $M_1 = 5$ and $M_2 = 1$ (22% of all observations). Obviously at least one of the divisions is lying and the safe outcome ($T_1 = T_2 = 3$) is the best response to diametrically opposed messages. Nevertheless, 48% of CMs follow the message from one division and ignore the other ($T_1 = T_2 = 1$ or $T_1 = T_2 = 5$). This type of error reduces CMs' average payoffs (67.1 vs 61.2 ECUs), but they don't learn to stop making this mistake (48% in Rounds 1 – 9 vs. 49% in Rounds 10 – 18).

Errors in using the information contained in divisions' messages explain why CMs in the Centralization treatment do no better than a babbling equilibrium. Suppose they had adopted the following "simple" strategy. If the divisions' messages agree, choose the product type that matches their messages. If $M_1 = 2$ and $M_2 = 1$, set $T_1 = T_2 = 2$. If $M_1 = 5$ and $M_2 = 4$, set $T_1 = T_2 = 4$. The preceding two cases cover cases where the divisions "almost" agree. Choose the safe outcome ($T_1 = T_2 = 3$) for all other message pairs. This simple strategy avoids the most obvious errors and would have yielded an average total surplus of 72.6 across the three treatments, which is significantly more than either the 69.4 that CMs actually achieved ($n = 59$; $z = 5.42$; $p < .01$) or the babbling equilibrium ($n = 59$; $z = 6.47$; $p < .01$).

Result 9: *CMs systematically make errors using the information contained in divisions' messages. These mistakes are responsible for their failure to beat the babbling equilibrium.*

5.2.b. A Failure To Communicate: Contrary to our hypotheses, neither treatment combining decentralization with structured communication (Division Suggestions and Manager Suggestions) has much impact on total surplus. This subsection examines the messages sent in these two treatments to better understand the lack of improvement relative to Baseline-Decentralized.

The Division Suggestions treatment creates a different but still challenging coordination problem. Rather than having to coordinate on product types, divisions have to coordinate on messages. If divisions agree on a product type by the third round of communication, they almost always coordinate their product choices (96%). The problem is that the divisions only reach an agreement in 63% of the observations.

[17] This refers to an empirical best response, not a theoretical construction. For each combination of messages observed in the data, we calculate the average payoff the CM would have achieved by choosing $T_1 = T_2 = 1$, $T_1 = T_2 = 2$, etc. The best response is the product type that maximizes the CM's average payoff.

Experience does not help much; the agreement rate is only 66% in Rounds 10 – 18. The divisions sometimes coordinate without an agreement, but the rate is low (40%).

Making matters worse, agreements do not make coordination more likely to be efficient. Subject to coordinating, the rate of efficient coordination is 53% with an agreement versus 49% without. This follows from the nature of agreements. For games where the efficient and safe equilibria do not coincide ($G \neq 3$), 62% of agreements call for play of the safe equilibrium. When the divisions agree on the efficient equilibrium (32% of agreements when $G \neq 3$), they always play the agreed upon equilibrium. The problem is that divisions usually do not agree on an efficient equilibrium. This mirrors the problem we observed in Baseline-Decentralization – the safe equilibrium provides a straightforward way of solving the difficult coordination problem faced by divisions.

Turning to the Manager Suggestions treatment, two things have to happen for CMs to exercise effective leadership: CMs have to make useful suggestions and divisions need to follow them. In practice, neither step reliably occurs.

Table 4 summarizes the suggestions made by CMs as a function of the game. We pool Games 1 and 5 and Games 2 and 4, which are mirror images of each other. When suggestions of the safe and efficient equilibria coincide in Game 3, we classify the CM as having suggested the safe equilibrium. CMs often don't give good advice. It is striking that 15% of messages do not even call for the divisions to coordinate their product types. Given that CMs in the centralization treatments almost always choose the same product type for the two divisions, it is hard to fathom why CMs would give advice to *not* coordinate. Even when coordination is advised, CMs often take a conservative approach by calling for coordination at the safe equilibrium. Excluding Game 3, where the safe and efficient equilibria coincide, only 36% of all messages suggest the efficient equilibrium. This percentage increases with experience, but only slightly, rising from 34% in Rounds 1 – 9 to 38% in Rounds 10 – 18.

Table 4: Manager Suggestions by Game Class

	Games 1 and 5	Games 2 and 4	Game 3
Non-equilibrium	16.4%	15.9%	13.2%
Safe Equilibrium	41.2%	32.8%	78.1%
Efficient Equilibrium	32.2%	40.5%	n/a
Other Equilibrium	10.2%	10.8%	8.8%

Note: When suggestions of the safe and efficient equilibria coincide in G = 3, we classify the CM as having suggested the safe equilibrium.

The conservative approach of CMs makes sense if suggesting the safe equilibrium leads to coordination when none would occur otherwise, but this is not the case. To set terminology, we define the

CM's "realized suggestion" as the suggestion made *for the realized game*. The probability that the divisions coordinate is high when the realized suggestion calls for either the safe or the efficient equilibrium (71% and 68% respectively, averaging across all games and all rounds). These figures compare well with 40% for all other realized suggestions as well as the 58% coordination rate in Baseline-Decentralization. It appears that realized suggestions of either the safe or efficient equilibrium generate higher coordination rates, somewhat justifying the conservative approach of many CMs. However, correlation does not imply causation. Table 4 shows that manager suggestions vary by game, and it is easier to coordinate in Game 3 than the other games. Additionally, given the major role of precedence in coordination games, both the suggestions and the outcomes may reflect lagged outcomes. For example, a CM may suggest the safe equilibrium because she knows the divisions have previously coordinated on the safe equilibrium. Thus, the association between realized suggestions of the safe equilibrium and coordination may reflect the circumstances when such suggestions are made rather than a causal link.

The regression in Table 5 is designed to establish a causal relationship between realized suggestions and outcomes. The data are all observations from the Manager Suggestions and Baseline-Decentralization treatments. The dependent variable is a dummy for coordinating ($T_1 = T_2$). Since this is a binary outcome, the regression uses a probit model. We report marginal effects with standard errors corrected for clustering at the group level.

Table 5: Regression Analysis, Effects Realized of Manager Suggestions on Coordination

Dependent Variable	Parameter Estimate	Standard Error
Realized Suggestion: Safe Equilibrium	0.064	0.075
Realized Suggestion: Efficient Equilibrium (Game ≠ 3)	0.175**	0.067
Realized Suggestion: Other Equilibrium	-0.223**	0.106
Realized Suggestion: Non-equilibrium	-0.096	0.109
Lagged Safe Equilibrium	0.260***	0.052
Lagged Efficient Equilibrium	0.110**	0.048
Lagged Other Equilibrium	0.319***	0.055

Note: The dataset contains 918 observations from 54 groups (clusters). Data from Round 1 are dropped to allow use of lagged variables. Standard errors are corrected for clustering at the group level. Three (***), two (**), and one (*) stars indicate statistical significance at the 1%, 5%, and 10% levels respectively.

The independent variables of primary interest are dummies for the CM's realized suggestion in the current round. These four variables (Realized Suggestion: Safe Equilibrium; Realized Suggestion: Efficient Equilibrium; Realized Suggestion: Other Equilibrium; Realized Suggestion: Non-equilibrium) capture all possible realized suggestions. The associated parameter estimates measure differences from Baseline-Decentralization. For example, the parameter estimate for "Realized Suggestion: Safe Equilibrium" captures how much higher the probability of coordinating is following a realized suggestion of the safe equilibrium in Manager Suggestions as compared to Baseline-Decentralization where no suggestions are possible. To separate the effect of suggestions from the effect of precedence, dummies for the possible lagged outcomes are included as independent variables (the omitted category is coordination failure). Data from the first round are dropped due to the use of a lagged variable. The model includes game and round dummies which are not reported to save space.

Controlling for the realized game and the lagged outcome, a realized suggestion of the safe equilibrium only has a small (and not significant) effect on the likelihood of coordinating relative to Baseline-Decentralization. The regression indicates that there is *not* a causal relationship between realized suggestions of the safe equilibrium and coordination. Realized suggestions of the efficient equilibrium do have a large and significant effect, indicating that a causal link does exist in this case.

If the divisions coordinate, they generally follow the CM's realized suggestion. When the realized suggestion is the safe equilibrium *and the divisions coordinate*, they almost always coordinate on the safe equilibrium (99%). Likewise, the divisions usually coordinate on the efficient equilibrium (92%) if it is the realized suggestion *and the divisions coordinate*. Analysis in Appendix D (Table D3) paralleling the regression in Table 5 implies that the primary effect of realized suggestions of the safe equilibrium is to shift play away from the efficient equilibrium towards the safe equilibrium. Allowing manager suggestions would be helpful if CMs primarily suggested play of the efficient equilibrium or if divisions responded to suggestions of the safe equilibrium by coordinating when they otherwise would not, but neither of these desirable outcomes occurs.

Result 10: *Performance in the Division Suggestions treatment is limited by frequent failures to reach an agreement between the divisions. In the Manager Suggestions treatment, little effect is observed due to poor advice by CMs and a failure of conservative advice to overcome coordination failure.*

5.2.c: The Impact of Chat Content on Coordination: To evaluate the impact of specific message types in Chat-Decentralization and Chat-Centralization, we developed a systematic scheme for coding message content. The goal was to quantify communication that might be relevant for the play of the game, avoiding prejudgments about which sorts of messages were important. We employed the methods developed by Cooper and Kagel (2005). After reading a random sample of conversations, we developed a coding scheme. Two research assistants then independently coded the content of all chat conversations. No effort was made to force agreement among coders. For several categories (marked with asterisks on Table 6), the initial two coders had a Cohen's kappa of less than .5, indicating relatively low agreement. These categories were all recoded by a third coder who was given extensive training in an attempt to improve the quality of the coding. The research assistants were not informed about any hypotheses the co-authors had about the messages. They were told that their job was to simply capture what had been said without concern to the possible effects of what had been said. Coding was binary – a message line was coded as a 1 if it was deemed to contain the relevant category of content and zero otherwise. We had no requirement on the number of codings for a message line – a coder could check as many or few categories as he or she deemed appropriate. A number of the categories also had sub-categories. For example, the coding scheme has a category for suggesting what product types should be chosen and sub-categories for specific suggestions (e.g. suggesting play of the efficient equilibrium). A coder was free to check whatever sub-categories they deemed appropriate when the corresponding category was checked off. Our analysis of the coding uses averages across coders unless otherwise noted.

Table 6: Frequency of Coding Categories

Coding Category	Chat-Decentralization	Chat-Centralization
Any Suggestion	93.1%	90.7%
Suggest Safe Equilibrium	54.1%	60.6%
Suggest Efficient Equilibrium	48.4%	57.7%
Agreement to Suggestion	78.9%	67.9%
Discuss Need to Coordinate *	6.4%	4.0%
Discuss Fairness *	31.8%	43.9%
Discuss Efficiency	39.4%	37.7%
Questions About Rules of the Experiment *	11.7%	15.0%
Questions About How to Play *	10.9%	14.2%
Explanation *	21.7%	32.3%
Ask What Game Is Being Played (CM)	n/a	19.4%
Truthfully Reveal Game	n/a	68.4%
Lie About Game	n/a	3.4%
Contradict (One tells truth, other lies)	n/a	2.5%

Table 6 reports the frequency of the coding categories across all rounds, broken down by treatment. Statistics in this subsection are based on all rounds unless otherwise noted. Some of the categories are only relevant in Chat-Centralization. “Contradict” is not a category per se, but instead is a combination of the preceding two categories that accounts for cases where one division truthfully reported what game was being played and the other division lied. Table E1 in Appendix E provides a more complete description of the categories. The unit of observation is the entire conversation prior to play in a single round rather than a single message line within that conversation or messages from only one individual in the conversation.

So, for example, in 93.1% of the pre-play dialogues in Chat-Decentralization, at least one of the divisions suggested what product types should be chosen.

Starting with Chat-Decentralization, recall that this treatment significantly improves total surplus relative to Baseline-Decentralization while the parallel treatment with structured communication, Division Suggestions, does not.[18] Performance in Division Suggestions is limited by failures to reach an agreement as well as a strong tendency to agree to play the safe equilibrium rather than the efficient equilibrium. Agreements are more frequent in Chat-Decentralization (63% agreement in Division Suggestions vs. 79% in Chat-Decentralization). Given that divisions almost always coordinate if an agreement is reached (95%), this translates directly into improved coordination and, by extension, improved performance. Chat-Decentralization does not solve the second problem that plagued the Division Suggestions treatment. In cases where the safe and efficient equilibria do not coincide ($G \neq 3$), only 34% of agreements call for play of the efficient equilibrium. When the divisions agree to play the efficient equilibrium, they almost always do so (96%). Chat-Decentralization fails to do better than either Baseline-Centralization or the babbling equilibrium because such agreements occur too rarely. The divisions fail to use their information.

Turning to Chat-Centralization, the biggest problem facing a CM is determining what game is being played. CMs in Chat-Centralization get remarkably good information about this. In 68% of the observations, at least one division truthfully reveals the game being played. Only in 3% of observations does a division lie. These figures improve with experience from 65% and 5% in Rounds 1 – 9 to 72% and 2% in Rounds 10 - 18. Information transmission is far more efficient in Chat-Centralization than Baseline-Centralization. In Baseline-Centralization, CMs receive conflicting reports in 69% of the observations.[19] They get information, but it is difficult to interpret and, as documented in Section 5.2.a, the CMs struggle with extracting information from the divisions' messages. In Chat-Centralization, extracting information is simple. The CM receives some report about the game in 69% of the observations. For 95% of these cases, the CM receives a truthful report without contradiction.[20] Almost always, the CM either has no information and doesn't face an information extraction problem or has unambiguous information that makes information extraction trivial.

The high quality of information transmission is enormously important for efficiency. In games where the safe and efficient equilibria do not coincide ($G \neq 3$), the frequency of efficient coordination rises from 18% when neither division truthfully reveals the game to 52% when at least one branch tells the

[18] Division Suggestions is parallel to Chat-Decentralization because both involve communication between divisions. Manager Suggestions is not parallel since communication is from CMs to divisions rather than between divisions.

[19] In these cases, usually (80%) one division is telling the truth while the other is lying.

[20] While rare (1%), there are also a few cases where the CM receives a false report without contradiction.

truth.[21] The frequency of efficient coordination is almost unaffected when one division tells the truth and the other division lies (53%), albeit based on a very small number of observations. Unlike politics, the truth wins in this environment.

The nature of truth-telling is quite different in Chat-Centralization than in Baseline-Centralization. In Baseline-Centralization, most subjects mix between telling the truth and lying. 69% of the divisions both tell the truth in at least a third of the rounds and lie in at least a third of the rounds. There are only two divisions that never lie and none that never tell the truth. As noted previously, partial lying is common. Divisions are strategic when choosing to tell the truth, doing so more often when it is to their benefit to be believed. This can be seen in Table 7. The games have been remapped for the D2 role so all observations are from the point of view of D1 (i.e. Equilibrium 1 is the worst equilibrium and Equilibrium 5 is the best).[22] Divisions are least likely to tell the truth and most likely to tell a lie when the efficient equilibrium would be worst for them (Game 1). This flips when the efficient equilibrium would be best for them (Game 5).

Table 7: Frequency of Truth-Telling and Lying

Game (Remapped)	Chat-Centralization		Baseline-Centralization	
	Truth	Lie	Truth	Lie
1	46.6%	2.5%	24.9%	75.1%
2	42.8%	3.1%	36.4%	63.6%
3	43.0%	2.2%	68.0%	32.0%
4	50.1%	1.9%	73.3%	26.7%
5	47.1%	0.3%	87.6%	12.4%

These patterns change in Chat-Centralization. Mixing between truth-telling and lying is largely non-existent. There are 47 subjects in the division role who send at least one message reporting what game is being played. These 47 divisions send an average of 9.8 reports over the course of 18 rounds. 35 of 47 reporting divisions never lie and another 6 of 47 only lie once in 18 rounds. There are no subjects who lie in at least a third of the rounds and none who lie in more than 40% of their reports. By extension, there are no divisions that both tell the truth in at least a third of the rounds and lie in at least a third of the rounds.[23]

[21] It may be surprising that the coordination rate is not close to zero when there isn't a truthful report and $G \neq 3$. In 87% of these cases, there is a suggestion that the efficient equilibrium should be played. A possible interpretation is that these suggestions serve as an indirect method of revealing the game, making a direct report unnecessary.

[22] Remapping the games allows us to pool data for the two divisions into a single table. Game 1 for a D2 is mapped to Game 5, Game 2 for a D2 is mapped to Game 4, etc.

[23] There are three divisions who both tell the truth in at least a third of their reports (rather than rounds) and lie in at least a third of their reports.

Subjects mix, but it is almost entirely between telling the truth and not reporting. Subject to lying, partial lies are common (40% of lies), but in absolute terms partial lies are necessarily rare given the low overall rate of lying. Returning to Table 7, truth-telling is not sensitive to incentives. (Note that Table 7 reports the frequency that an *individual* division reports truthfully at some point during the pre-play communication. This is different from the figure reported in Table 6, which shows the frequency that *at least one* of the two divisions reports truthfully at some point during the pre-play communication.) Telling the truth is just about as likely when it is most advantageous for a division to lie (46.6% in Game 1) as when it is most advantageous to tell the truth (47.1% in Game 5).

One possible reason for infrequent lying is that divisions feel guiltier about lying to the CM when they have been directly asked for a report. However, it is surprisingly rare for the CM to request reports about what game is being played (19%). Requests for reports are somewhat more common in the early rounds (33% in Rounds 1 – 6) but fall off steadily over time (17% in Rounds 7 – 12 and 8% in Rounds 13 – 18). Asking for reports works, in the sense of increasing the frequency of reports about what game is being played (94% with a request vs. 63% without). Unfortunately, the quality of the reports is harmed. When the CM requests a report, the fraction of lies rises from 2% to 10%.

Another possibility is that divisions avoid lying because they are concerned about being "fact-checked." In both Baseline-Centralization and Chat-Centralization, the CM knows *ex post* when a division has lied, but in Chat-Centralization it is possible for the other division to call out a liar in real time. This is a valid fear since in 40% of the observations where a division lies, the other division corrects them.[24] It may be more embarrassing to be actively called out as a liar than to merely be revealed as a liar.[25]

Most non-reports arise in cases where divisions have good reason to consider a report unnecessary. After one division has revealed what game is being played there is little need for the other to do so (assuming they plan to report truthfully). In line with this, for 44% of the cases where a division does not make a report, the other division has reported truthfully. It is also pointless to report what game is being played if the safe equilibrium is going to be played regardless. Consider observations where the safe and efficient equilibria could be distinguished in the *previous* round ($G \neq 3$). If the safe equilibrium was played in the previous round, *neither* division reports what game is being played in the current round in 44% of the observations; if divisions expect the safe equilibrium to be played they tend to abstain from reporting since they expect that the central management will make a decision without taking into account any new

[24] This is different from the figure reported as "Contradict" in Table 6, which measures cases where one division reported truthfully and the other lied. The 40% figure refers to "fact-checking" where one division explicitly corrects a false report by the other. (e.g. "It is Game 3." "No, it is really Game 2.")

[25] Fact-checking partially explains why the frequency of the efficient equilibrium is the same when one division tells the truth and is not contradicted versus when one division tells the truth and the other contradicts them by lying, since fact-checking often gives the CM guidance about which division to believe.

information. By contrast, if they expect the efficient equilibrium to be played then they have an incentive to report to guide the manager's decisions. Indeed, when the efficient equilibrium was played in the previous round, it only occurs in 17% of the observations that neither division reports. In combination, 78% of non-reports occur in cases where either the other division has told the truth or the safe equilibrium is played. Most non-reports can be explained by a lack of any real need to report the game being played.

Table 8: Probit Regressions, Effects of Chat on Play

Treatment	Chat-Decentralization		Chat-Centralization
Dependent Variable	Coordination	Efficient Coordination	Efficient Coordination
Agreement	0.106*** (0.021)	0.217*** (0.060)	0.051 (0.078)
Discuss Need to Coordinate	-0.031 (0.046)	0.231 (0.180)	-0.157 (0.199)
Discuss Fairness	0.010 (0.023)	-0.264*** (0.076)	-0.107 (0.105)
Discuss Efficiency	0.008 (0.026)	0.272*** (0.063)	0.111 (0.090)
Questions About Rules	-0.053* (0.030)	-0.045 (0.148)	-0.168 (0.104)
Questions About Play	0.036 (0.033)	-0.130 (0.148)	-0.099 (0.149)
Explanation	-0.039 (0.031)	0.061 (0.125)	0.012 (0.098)
Ask What Game			0.127 (0.081)
Truthfully Reveal Game			0.294*** (0.094)
Contradict			-0.045 (0.234)

Notes: All models include 459 observations. Marginal effects are reported. In parentheses, we report standard errors corrected for clustering at the group level. All regressions include controls for the game being played, three-round block, and lagged outcomes. Coefficients for these variables are not reported to save space. Three (***), two (**), and one (*) stars indicate significance at the 1%, 5%, and 10% levels using two-tailed tests.

None of the preceding establishes a causal relationship between the content of pre-play communication and outcomes. Establishing causality is tricky because outcomes and the content of communication may both depend on lagged outcomes. Table 8 shows the results of probit regressions designed to establish causality. The left panel shows results for Chat-Decentralization and the right panel shows results for Chat-Centralization. The dependent variable is either a dummy for coordination (the two divisions' product types match) or efficient coordination. First round data are dropped to allow the use of lagged variables. There is not a regression for coordination in the Chat-Centralization treatment because there was 100% coordination following Round 1.

As independent variables, all regressions include the lagged outcome (dummies for coordination failure (Chat-Decentralization only), safe coordination, and efficient coordination), game dummies, and a dummy for late rounds (Rounds 10 – 18). These are not reported to save space in the table. All regressions include the average coding for the categories reported in Table 6 with the following exceptions. The categories for "Lie About Game" and "Contradict" are highly collinear, so we only include the latter (we felt this was the more interesting of the two). Including suggestions about what actions to play makes the regressions circular (subjects do what they say they should do), so these categories are omitted. We report marginal effects. Standard errors are corrected for clustering at the group level.

We previously mentioned the importance of agreements for coordination in Chat-Decentralization, and the regressions provide additional evidence of this. There is a strong positive relationship between reaching an agreement and either coordinating in general or specifically coordinating at the efficient equilibrium. Not surprisingly, discussing efficiency moves play towards the efficient equilibrium while discussing fairness has the opposite effect. Questions relating to the rules of the games are (weakly) negatively related to coordination. This suggests that coordination failure stems in part from confusion.

In Chat-Centralization, we previously noted that the efficient equilibrium is far more likely when at least one division reports truthfully. The regression supports this conclusion as the efficient equilibrium is significantly more likely to occur when at least one division truthfully reports what game is being played. Asking what game is being played also increases play of the efficient equilibrium, but the effect is weak and (barely) misses statistical significance ($p = .11$). Presumably the effect of receiving more reports after a request is somewhat offset by the lower quality of reports as described above. The combination of a true report and a lie ("Contradict") has little effect on outcomes, consistent with our comments above. Looking at the other categories, the only one that comes close to being significant ($p = .11$) is "Questions About Rules," suggesting that confusion plays a role in limiting play of the efficient equilibrium. It is interesting that, unlike Chat-Decentralization, discussions of fairness and efficiency do *not* have significant effects on whether the efficient equilibrium is played. Truthful reporting is the sole factor that serves as a strong predictor of the efficient equilibrium, highlighting the role of information transmission.

Result 11: *Chat leads to much better transmission of the divisions' information than occurs in Baseline-Centralization. This happens because divisions frequently tell the truth, almost never lie, and rarely confront CMs with conflicting reports. Truth-telling is strongly associated with play of the efficient equilibrium. The frequency of truth-telling, lying, and non-report in Chat-Centralization is not sensitive to what game is being played. Unlike Division Suggestions, the patterns of truth-telling in Chat-Centralization do not parallel what is typically reported in the literature on lying.*

6. Concluding Remarks: The question of whether to centralize or decentralize decision making comes up in numerous economic settings. Our experiments use a new game, the decentralization game, to study how organizations perform under centralization and decentralization. This game is designed to confront subjects with a challenging environment that accentuates differences between these two organizational structures. Ideally, the organization would coordinate on a single product type for both divisions *and* make use of the divisions' private information. *Ex ante*, there were good reasons to believe that both organizational structures would fail to reach these goals in the absence of pre-play communication. In most organizational settings, it is natural that communication occurs beyond bare reports of basic information about the state of the world. We therefore study how pre-play communication facilitates efficient coordination.

Without pre-play communication, total surplus is significantly higher under centralization than decentralization. This result is robust to changing the matching protocol (partners vs. strangers) or changing the payoffs to reduce the conflict between the divisions' interests. The superior performance under centralization reflects frequent coordination failure under decentralization. Performance under both organizational structures is limited by a strong tendency to coordinate on playing the safe equilibrium regardless of the state of the world. In other words, there is a general failure to take advantage of the divisions' information.

Structured communication in terms of either pre-play messages between divisions (Division Suggestions) or pre-play advice from the central manager (Manager Suggestions) does *not* lead to significantly higher total surplus than the baseline decentralization treatment without pre-play communication (Baseline-Decentralization). In Division Suggestions, performance is limited by the divisions' frequent inability to reach agreement and overuse of the safe equilibrium. Good advice from the central manager is quite effective in Manager Suggestions, but, unfortunately, central managers often fail to give good advice.

Free-form communication leads to higher surplus under either decentralization or centralization. However, the total surplus with chat under *decentralization* is no higher than in the parallel baseline treatment under *centralization*, nor does it beat the predicted total surplus from the babbling equilibrium. In contrast, the combination of chat with centralization yields significantly higher surplus than the baseline under centralization, chat under decentralization, and the babbling equilibrium. The combination of chat and centralization succeeds in gaining approximately half of the possible gains over the babbling equilibrium. Underlying this strong performance is significantly higher use of the efficient equilibrium. Only when chat and centralization are combined is the group able to achieve coordination *and* make effective use of the divisions' information. At times communication, especially free-form chat, can seem like the magic bullet in experimental economics, solving all problems. This is not quite the case in our experiments. Chat always improves matters, but the choice of an organizational structure remains important in determining how the organization functions.

Combining chat and centralization leads to more use of the efficient equilibrium because of the high quality of information transmission. In the baseline treatment with centralization, divisions' reports look much like what would be expected from the literature studying lies and lie aversion. Divisions mix between telling the truth and lying, the frequency of truth-telling responds strongly to incentives, and partial lies are common. These expected patterns vanish in the Chat-Centralization treatment. Divisions almost never lie and the frequency of truth-telling does not respond to incentives. In Baseline-Centralization, the modal outcome is that the central manager receives conflicting reports and is faced with a difficult information extraction problem. Central managers do poorly at extracting what information is available in the reports. No such problem exists in Chat-Centralization. In almost all observations, the central manager either only receives truthful reports or receives no reports at all, but does not receive conflicting reports. When she receives accurate information, adopting the efficient equilibrium becomes a far simpler task.

The patterns of communication about the state of the world (i.e. what game is being played) in Chat-Centralization are quite different from what is typically observed in experiments on truth-telling. This could reflect a difference between chat and structured communication, but the results of Lundquist et al. suggest otherwise since they frequently observe deceptive lies (40%) with free-form messages. We speculate that having two identically informed divisions plays a major role in this finding by making it possible that a lie will be contradicted in real time. It is one thing to have somebody find out after the fact that you have lied, but it is presumably much more embarrassing to be called out mid-conversation. In the rare cases where lies are told, we often (40%) see fact-checking of this type.

In Chat-Centralization the central manager and the two divisions can freely discuss the available options with all three being on equal footing. All three parties can send messages, which are seen by all three. The political science literature refers to this type of open multi-lateral chat as "deliberation."[26] It has been suggested (Dawes et al.; 1990, Orbell et al., 1988; Dryzek and List, 2003) that the process of deliberation makes group members more willing to take into account the interests of the whole group (perhaps by increasing group identity). Along similar lines, the reluctance to lie in Chat-Centralization may in part be a result of divisions taking into account the interests of the organization as a whole.

Centralization works far better than decentralization in our experimental environment, but it would be obviously excessive to claim this is a universal result. The game we study is intentionally simple and

[26] See Myers and Mendelberg (2013) and Karpowitz and Mendelberg (2011) for overviews of research on political deliberation and of experimental work on the topic, respectively.

designed to accentuate the tradeoffs between the two organizational forms. In the future we plan to extend our study by adding complexities, such as group decision making, that mirror real-world organizations.

REFERENCES

Abeler, Johannes, Daniele Nosenzo, and Collin Raymond (forthcoming), "Preferences for Truth-Telling," *Econometrica*.

Alonso, Ricardo, Wouter Dessein, and Niko Matouschek (2008a), "When Does Coordination Require Centralization?" *American Economic Review*, 98(1), 145-179.

Alonso, Ricardo, Wouter Dessein, and Niko Matouschek (2008b), "Centralization versus Decentralization: An Application to Price Setting by a Multi-Market Firm," *Journal of the European Economic Association*, 6 (2-3), 457-467.

Alonso, Ricardo, Wouter Dessein, and Niko Matouschek (2015), "Organizing to Adapt and Compete," *American Economic Journal: Microeconomics*, 7(2), 158-187.

Aumann, Robert (1990), "Nash Equilibria are not Self-Enforcing," in J. J. Gabszewicz, J. F. Richard, and L. A. Wolsey (eds.), *Economic Decision-Making: Games, Econometrics and Optimisation: Contributions in Honour of Jacques H. Drèze*, 201–206. Amsterdam: Elsevier.

Battaglini, Marco (2002), "Multiple Referrals and Multidimensional Cheap Talk," *Econometrica*, 70(4), 1379–1401.

Ben-Ner, Avner, and Louis Putterman (2009), "Trust, communication and contracts: An experiment," *Journal of Economic Behavior & Organization*, 70(1-2), 106-121.

Blume, Andreas and Andreas Ortmann (2007), "The Effects of Costless Pre-play Communication: Experimental Evidence from Games with Pareto-ranked Equilibria," *Journal of Economic Theory*, 132, 274–290.

Blume, Andreas, Peter H. Kriss, and Roberto A. Weber (2017), "Pre-play communication with forgone costly messages: experimental evidence on forward induction," *Experimental Economics*, 20(2), 368-395.

Bolton, Gary, Jordi Brandts, and Axel Ockenfels (2005), "Fair Procedures: Evidence from Games Involving Lotteries," *Economic Journal*, 115, 1054-1076.

Brandts, Jordi, Gary Charness, and Matthew Ellman (2016), "Let's Talk: How Communication Affects Contract Design," *Journal of the European Economic Association,* 14(4), 943-974.

Brandts, Jordi and David J. Cooper (2007), "It's What You Say Not What You Pay: An Experimental Study of Manager-employee Relationships in Overcoming Coordination Failure," *Journal of the European Economic Association*, 5(6), 1223-1268.

Brandts, Jordi, David J. Cooper, and Roberto A. Weber (2015), "Legitimacy, Communication, and Leadership in the Turnaround Game," *Management Science*, 61(11), 2627-2645.

Cai, Hongbin, and Joseph Tao-Yi Wang (2006), "Overcommunication in strategic information transmission games," *Games and Economic Behavior*, 56(1), 7-36.

Camerer, Colin F. (2003), *Behavioral Game Theory*. Princeton University Press, Princeton, New Jersey.

Camerer, Colin F. and Roberto A. Weber. (2012). "Experimental Organizational Economics." In Robert Gibbons and John Roberts (Eds.), *Handbook of Organizational Economics*. Princeton University Press, Princeton, New Jersey.

Cason, Timothy N., and Vai-Lam Mui (2007), "Communication and coordination in the laboratory collective resistance game," *Experimental Economics*, 10(3), 251-267.

Charness, Gary and Martin Dufwenberg (2006), “Promises and Partnership,” *Econometrica*, 74, 1579-1601.

Charness, Gary, and Martin Dufwenberg (2010), “Bare promises: An experiment,” *Economics Letters*, 107(2), 281-283.

Cooper, David J., and John H. Kagel (2005), “Are two heads better than one? Team versus individual play in signaling games,” *American Economic Review,* 95(3), 477-509.

Cooper, David and Kai-Uwe Kühn (2014), “Communication, Renegotiation, and the Scope for Collusion,” *American Economic Journal: Microeconomics*, 6, 247-278.

Cooper, David and Roberto Weber (forthcoming), “Recent Advances in Experimental Coordination Games,” in M. Capra, R. Croson, M. Rigdon, and T. Rosenblatt (eds.), *The Handbook of Experimental Economic Game Theory*, Cheltenham: Edward Elgar Publishing.

Cooper, Russell, Douglas DeJong, Robert Forsythe, and Thomas Ross (1989), “Communication in the Battle of the Sexes Game: Some Experimental Results,” *RAND Journal of Economics,* 20, 568-587.

Dawes, Robyn M., Jeanne McTavish, and Harriet Shaklee (1977), "Behavior, communication, and assumptions about other people's behavior in a commons dilemma situation," *Journal of Personality and Social Psychology*, 35(1), 1.

Dawes, Robyn M., Alphons J.C. van de Kragt, and John M. Orbell (1990), “Cooperation for the benefit of us—Not me, or my conscience,” in J. J. Mansbridge (ed.), *Beyond self-interest* (pp. 97-110), Chicago: University of Chicago Press.

Dessein, Wouter, Luis Garicano and Robert Gertner (2010), “Organizing for Synergies,” *American Economic Journal: Microeconomics*, 2, 77-114.

Dryzek, John and Christian List (2003), “Social Choice Theory and Deliberative Democracy: A Reconciliation,” *British Journal of Political Science*, 33(1), 1-28.

Dufwenberg, Martin and Martin A. Dufwenberg (2018), “Lies in Disguise: A Theoretical Analysis of Cheating,” *Journal of Economic Theory*, 175, 248-264.

Erat, Sanjiv and Uri Gneezy (2012), “White Lies,” *Management Science*, 58(4) 723-733.

Evdokimov, Piotr and Umberto Garfagnini (2018), “Communication and Behavior in Organizations: An Experiment,” Working Paper.

Feri, Francesco, Bernd Irlenbusch, and Matthias Sutter (2010), “Efficiency Gains from Team-based Coordination – Large-scale Experimental Evidence,” *American Economic Review*, 100(4), 1892-1912.

Fischbacher, Urs and Franziska Föllmi-Heusi (2013), “Lies in Disguise – An Experimental Study on Cheating,” *Journal of the European Economic Association*, 11(3):525–547.

Fréchette, Guillaume R. (2015), “Laboratory Experiments: Professionals versus Students”, in G. Fréchette and A. Schotter (eds.), *Handbook of Experimental Economic Methodology*, Oxford: Oxford University Press.

Gintis, Herbert (2014), The Moral Basis of the Business World. Working paper.

Gneezy, Uri (2005), “Deception: The Role of Consequences,” *American Economic Review*, 95(1), 384-394.

Gneezy, Uri, Agne Kajackaite, and Joel Sobel (2018), “Lying Aversion and the Size of the Lie,” *American Economic Review*, 108(2), 419-453.

Hamman, John and Miguel A. Martínez-Carrasco (2018), “Delegation and Team Selection in Organizations: An Experimental Study,” Working Paper.

Hart, Oliver and John Moore (2008), “Contracts as Reference Points,” *Quarterly Journal of Economics*, 123, 1-48.

Hart, Oliver and Bengt Holmstrom (2010), “A Theory of Firm Scope,” *Quarterly Journal of Economics*, *125*(2), 483-513.

Hayek, Friedrich von (1945), “The Use of Knowledge in Society,” *American Economic Review,* 35(4), 519 – 530.

Isaac, R. Mark, and James M. Walker (1988), "Communication and free-riding behavior: The voluntary contribution mechanism," *Economic Inquiry*, 26(4), 585-608.

Karpowitz, Christopher and Tali Mendelberg (2011), “An Experimental Approach to Citizen Deliberation,” in Druckman, J. N., D. P. Green, J. H. Kuklinski, and A. Lupia (eds), *Cambridge Handbook of Experimental Political Science*, Cambridge: Cambridge University Press.

Kriss, Peter H., Andreas Blume, and Roberto A. Weber (2016), “Coordination with decentralized costly communication,” *Journal of Economic Behavior & Organization*, 130, 225-241.

Kriss, Peter H. and Roberto A. Weber. (2013), “Organizational formation and change: lessons from economic laboratory experiments.” In Anna Grandori (Ed.) *Handbook of Economic Organization: Integrating Economic and Organization Theory,* Cheltenham: Edward Elgar.

Kuzmics, Christoph, Thomas Palfrey, and Brian W. Rogers (2014), “Symmetric play in repeated allocation games,” *Journal of Economic Theory*, 154, 25–67.

Lange, Oskar and Fred M. Taylor (1938), *On The Economic Theory of Socialism,* Minneapolis: University of Minnesota Press.

Lundquist, Tobias, Tore Ellingsen, Erik Gribbe, and Magnus Johannesson (2009), “The aversion to lying,” *Journal of Economic Behavior & Organization*, 70(1-2), 81-92.

McElheran, Kristina (2014), “Delegation in Multi-Establishment Firms: The Organizational Structure of I.T. Purchasing Authority,” *Journal of Economics and Management Strategy*, 23(2), 225-257.

Mookherjee, Dilip (2006), “Decentralization, Hierarchies and Incentives: A Mechanism Design Perspective,” *Journal of Economic Literature*, 44(2), 367-390.

Myers, Daniel and Tali Mendelberg (2013), Oxford Handbook of Political Psychology. Oxford University Press, Ch. Political Deliberation, 699-734.

Orbell, John, Alphons J. C. van de Kragt, Robyn M. Dawes (1988), “Explaining Discussion-Induced Cooperation,” *Journal of Personality and Social Psychology*, 54(5), 811-819.

Ostrom, Elinor, James Walker, and Roy Gardner (1992), "Covenants with and without a sword: Self-governance is possible," *American Political Science Review*, 86(2), 404-417.

Rantakari, Heikki (2008), “Governing Adaptation,” *The Review of Economic Studies*, 75, 1257-1285.

Sánchez-Pagés, Santiago, and Marc Vorsatz (2007), “An experimental study of truth-telling in a sender–receiver game,” *Games and Economic Behavior*, 61(1), 86-112.

Thomas, Catherine (2011), “Too Many Products: Decentralized Decision Making in Multinational Firms,” *American Economic Journal: Microeconomics*, 3, 280-306.

Van Huyck, John, Ann Gillette, and Raymond Battalio (1992), “Credible assignments in coordination games,” *Games and Economic Behavior*, 4(4), 606-626.

Vespa, Emanuel and Alistair J. Wilson (2016), “Communication With Multiple Senders: An Experiment,” *Quantitative Economics*, 7(1), 1-36.

Note: Appendices are for online publication.

Appendix A: Proof of Theorem

Lemma: For any beliefs, the central manager (CM) will choose the same actions for the two divisions (D1 and D2).

Proof: Suppose not. This implies that the CM is choosing an outcome that is not a Nash equilibrium if the two divisions are allowed to choose their own actions. Either of the divisions could improve its payoff by switching to the action chosen by the other division. Moreover, the other division's payoff is also increased by this change. Since the CM's payoff equals the sum of the two divisions' payoffs, the CM's payoff also increases. This implies that the CM's initial choice could not have been optimal.

Given the preceding lemma, we can refer to the CM as choosing a single action in response to the divisions' messages.

Theorem: There does not exist a pure-strategy PBE where the CM chooses different actions for two different states of the world.

Proof: Suppose that such an equilibrium existed. Let S1 and S2 be two states where different actions are chosen. Let A1 and A2 be the actions chosen by the CM in equilibrium in S1 and S2 respectively, where A1 ≠ A2. Without loss of generality, assume that D1 prefers the outcome in S1 and D2 prefers the outcome in S2. Let M_i^j be the message sent by Di in Sj.

It cannot be the case that $M_1^1 = M_1^2$. Proof is by contradiction. Suppose $M_1^1 = M_1^2$. This implies that the CM's choice is determined solely by D2's message. Since D2 prefers A2, it should always send M_2^2 whether the true state of the world is S1 or S2. But then the CM would choose A2 in both S1 and S2. A contradiction follows. By the same logic, $M_2^1 \neq M_2^2$.

Suppose that D1 deviates by sending M_1^1 in S2. The resulting pair of messages (M_1^1, M_2^2) cannot make D1 better off than A2 or a profitable deviation from equilibrium exists. It follows that (M_1^1, M_2^2) leads to an outcome that makes D1 worse off (weakly) than A2. However, because the two divisions' preferences over possible outcomes are diametrically opposed, this implies that D2 can gain by sending M_2^2 in S1, giving D2 a profitable deviation from equilibrium. A contradiction follows. **Q.E.D.**

Appendix B: INSTRUCTIONS
DECENTRALIZATION TREATMENT

Thanks for coming to the experiment. You will receive 5 euros for participation in the experiment. Also, you will earn additional money during the experiment.

Participants have been randomly assigned to one of three roles: F, C and A. This role will be the same throughout the experiment.

There will be 18 separate periods. We will now present the instructions for the first block of nine periods. Later you will receive further instructions. In each period, you will be in a group of three participants, one in each role. The persons that you are matched with will change from period to period. During the nine periods you will never meet another person twice. Also, at no time will you know the identity of who you are matched with.

Each period is independent from the others and develops in the following way. At the beginning of the period, the computer will randomly determine which of the following five games will be played.

In each of the cells the first number shown in yellow is the payoff that the person in the F role will receive, the second number shown in green is the payoff that the person in the C role will receive and the third number shown in red is the payoff for the person in the A role. As you can see all five games have five rows: f1, f2, f3, f4 and f5 [Note: The Spanish word for row is “fila”. We have kept the original abbreviations in the payoff tables.], and five columns; c1, c2, c3, c4 and c5. [Note: The Spanish word for row is “fila”. We have kept the original abbreviations in the payoff tables shown below.] Observe also that the numbers in the different cells differ between the games.

Game 1

	c1	c2	c3	c4	c5
f1	26, 54, 80	12, 29, 41	-2, 4, 2	-16, -21,-37	-30,-46,-76
f2	15, 40, 55	29, 43, 72	15, 18, 33	1, -7, -6	-13, -32,-45
f3	4, 26, 30	18, 29, 47	32,32, 64	18, 7, 25	4, -18, -14
f4	-7, 12, 5	7, 15, 22	21, 18, 39	35, 21, 56	21, -4, 17
f5	-18, -2, -20	-4, 1,-3	10, 4, 14	24, 7, 31	38, 10, 48

Game 2

	c1	c2	c3	c4	c5
f1	22, 50, 72	8, 33, 41	-6, 8, 2	-20, -17, -37	-34,-42,-76
f2	19, 36, 55	33, 47, 80	19, 22, 41	5, -3, 2	-9, -28,-37
f3	8, 22, 30	22,33, 55	36, 36, 72	22, 11, 33	8, -14, -6
f4	-3,8, 5	11,19,30	25, 22, 47	39, 25, 64	25, 0, 25
f5	-14,-6,-20	0,5,5	14, 8, 22	28, 11, 39	42, 14, 56

Game 3

	c1	c2	c3	c4	c5
f1	18, 46, 64	4, 29, 33	-10,12, 2	-24, -13, -37	-38, -38, -76
f2	15, 32, 47	29, 43, 72	15, 26, 41	1, 1, 2	-13, -24,-37
f3	12, 18, 30	26, 29, 55	40, 40, 80	26, 15, 41	12, -10,2
f4	1,4, 5	15, 15, 30	29,26, 55	43,29, 72	29, 4, 33
f5	-10,-10,-20	4,1, 5	18,12, 30	32, 15, 47	46, 18, 64

Game 4

	c1	c2	c3	c4	c5
f1	14, 42, 56	0, 25, 25	-14, 8,-6	-28, -9, -37	-42, -34, -76
f2	11, 28, 39	25, 39, 64	11, 22, 33	-3, 5, 2	-17,-20,-37
f3	8, 14, 22	22, 25, 47	36, 36, 72	22, 19, 41	8, -6, 2
f4	5, 0, 5	19, 11, 30	33, 22, 55	47, 33, 80	33, 8, 41
f5	-6,-14,-20	8, -3, 5	22, 8, 30	36, 19, 55	50, 22, 72

Game 5

	c1	c2	c3	c4	c5
f1	10, 38, 48	-4, 21, 17	-18, 4,-14	-32,-13,-45	-46,-30, -76
f2	7, 24, 31	21, 35, 56	7, 18, 25	-7, 1,-6	-21,-16,-37
f3	4, 10, 14	18, 21, 39	32, 32, 64	18, 15, 33	4, -2, 2
f4	1,-4,-3	15, 7, 22	29, 18, 47	43, 29, 72	29, 12, 41
f5	-2,-18,-20	12,-7, 5	26,4, 30	40, 15, 55	54, 26, 80

Each of the five games has the same chance of being chosen in each period separately. That is in each period, each of the games will be chosen with 20%probability. Player F and player C will be informed of which game has been chosen, but player A will not be informed of which game has been chosen.

After having seen which game has been selected by the random draw, players F and player C will separately make decisions. Player F will choose between f1, f2, f3, f4 and f5 and player C will choose between columns c1, c2, c3, c4 and c5. Player A will not make any decisions.

The payoffs of players F, C and A will be the ones in the cell determined by the row chosen by F and the column chosen by C for the game selected by the random draw. Remember that players F and C will make their decisions independently from each other.

After each period everybody will be informed about what row was chosen by F and what column was chosen by C sent and about which game was randomly selected.

After this, a new period will start which will develop in the same way until reaching period 9. Remember that the persons you play with will change from period to period.

Each ECU is worth 0,02 euros. At the end of the session you will receive 5 euros plus what you will have earned in all 18 rounds of the experiment.

You can ask questions at any time. If you have a question, please raise your hand and one of us will come to your place to answer it.

INSTRUCTIONS

CENTRALIZATION TREATMENT

Thanks for coming to the experiment. You will receive 5 euros for participation in the experiment. Also, you will earn additional money during the experiment.

Participants have been randomly assigned to one of three roles: F, C and A. This role will be the same throughout the experiment.

There will be 18 separate periods. We will now present the instructions for the first block of nine periods. Later you will receive further instructions. In each period, you will be in a group of three participants, one in each role. The persons that you are matched with will change from period to period. During the nine periods you will never meet another person twice. Also, at no time will you know the identity of who you are matched with.

Each period is independent from the others and develops in the following way. At the beginning of the period, the computer will randomly determine which of the following five games will be played.

In each of the cells the first number shown in yellow is the payoff that the person in the F role will receive, the second number shown in green is the payoff that the person in the C role will receive and the third number shown in red is the payoff for the person in the A role. As you can see all five games have five rows: f1, f2, f3, f4 and f5, and five columns; c1, c2, c3, c4 and c5. Observe also that the numbers in the different cells differ between the games.

Game 1

	c1	c2	c3	c4	c5
f1	26, 54, 80	12, 29, 41	-2, 4, 2	-16, -21,-37	-30,-46,-76
f2	15, 40, 55	29, 43, 72	15, 18, 33	1, -7, -6	-13, -32,-45
f3	4, 26, 30	18, 29, 47	32,32, 64	18, 7, 25	4, -18, -14
f4	-7, 12, 5	7, 15, 22	21, 18, 39	35, 21, 56	21, -4, 17
f5	-18, -2, -20	-4, 1,-3	10, 4, 14	24, 7, 31	38, 10, 48

Game 2

	c1	c2	c3	c4	c5
f1	22, 50, 72	8, 33, 41	-6, 8, 2	-20, -17, -37	-34,-42,-76
f2	19, 36, 55	33, 47, 80	19, 22, 41	5, -3, 2	-9, -28,-37
f3	8, 22, 30	22,33, 55	36, 36, 72	22, 11, 33	8, -14, -6
f4	-3,8, 5	11,19,30	25, 22, 47	39, 25, 64	25, 0, 25
f5	-14,-6,-20	0,5,5	14, 8, 22	28, 11, 39	42, 14, 56

Game 3

	c1	c2	c3	c4	c5
f1	18, 46, 64	4, 29, 33	-10,12, 2	-24, -13, -37	-38, -38, -76
f2	15, 32, 47	29, 43, 72	15, 26, 41	1, 1, 2	-13, -24,-37
f3	12, 18, 30	26, 29, 55	40, 40, 80	26, 15, 41	12, -10,2
f4	1,4, 5	15, 15, 30	29,26, 55	43,29, 72	29, 4, 33
f5	-10,-10,-20	4,1, 5	18,12, 30	32, 15, 47	46, 18, 64

Game 4

	c1	c2	c3	c4	c5
f1	14, 42, 56	0, 25, 25	-14, 8,-6	-28, -9, -37	-42, -34, -76
f2	11, 28, 39	25, 39, 64	11, 22, 33	-3, 5, 2	-17,-20,-37
f3	8, 14, 22	22, 25, 47	36, 36, 72	22, 19, 41	8, -6, 2
f4	5, 0, 5	19, 11, 30	33, 22, 55	47, 33, 80	33, 8, 41
f5	-6,-14,-20	8, -3, 5	22, 8, 30	36, 19, 55	50, 22, 72

Game 5

	c1	c2	c3	c4	c5
f1	10, 38, 48	-4, 21, 17	-18, 4,-14	-32,-13,-45	-46,-30, -76
f2	7, 24, 31	21, 35, 56	7, 18, 25	-7, 1,-6	-21,-16,-37
f3	4, 10, 14	18, 21, 39	32, 32, 64	18, 15, 33	4, -2, 2
f4	1,-4,-3	15, 7, 22	29, 18, 47	43, 29, 72	29, 12, 41
f5	-2,-18,-20	12,-7, 5	26,4, 30	40, 15, 55	54, 26, 80

Each of the five games has the same chance of being chosen in each period separately. That is in each period, each of the games will be chosen with 20%probability. Player F and player C will be informed of which game has been chosen, but player A will not be informed of which game has been chosen.

After having seen which game has been selected by the random draw, players F and player C will separately send messages to player A saying which game has been selected . This message can be truthful or not. Once player A has received the messages he will choose a row and column without knowing which game was selected.

The payoffs of players F, C and A will be the ones in the cell determined by the row and the column chosen by A for the game selected by the random draw. Remember that players F and C willsend their messages independently from each other.

After each period everybody will be informed about what row and what column was chosen by A and about which game was randomly selected.

After this, a new period will start which will develop in the same way until reaching period 9. Remember that the persons you play with will change from period to period.

Each ECU is worth 0,02 euros. At the end of the session you will receive 5 euros plus what you will have earned in all 18 rounds of the experiment.

You can ask questions at any time. If you have a question, please raise your hand and one of us will come to your place to answer it.

Appendix C:
Table C1: Stage Game Payoffs (k1 = 54, k2 = 7, k3 = 4, and k4 = 14)

Note: Each cell contains the payoffs for D1 (π_{D1}), D2 (π_{D2}), and CM (π_{CM}).

Game 1

	C1	C2	C3	C4	C5
R1	26, 54, 80	12, 29, 41	-2, 4, 2	-16, -21,-37	-30,-46,-76
R2	15, 40, 55	29, 43, 72	15, 18, 33	1, -7, -6	-13, -32,-45
R3	4, 26, 30	18, 29, 47	32,32, 64	18, 7, 25	4, -18, -14
R4	-7, 12, 5	7, 15, 22	21, 18, 39	35, 21, 56	21, -4, 17
R5	-18, -2, -20	-4, 1,-3	10, 4, 14	24, 7, 31	38, 10, 48

Game 2

	C1	C2	C3	C4	C5
R1	22, 50, 72	8, 33, 41	-6, 8, 2	-20, -17, -37	-34,-42,-76
R2	19, 36, 55	33, 47, 80	19, 22, 41	5, -3, 2	-9, -28,-37
R3	8, 22, 30	22,33, 55	36, 36, 72	22, 11, 33	8, -14, -6
R4	-3,8, 5	11,19,30	25, 22, 47	39, 25, 64	25, 0, 25
R5	-14,-6,-20	0,5,5	14, 8, 22	28, 11, 39	42, 14, 56

Game 3

	C1	C2	C3	C4	C5
R1	18, 46, 64	4, 29, 33	-10,12, 2	-24, -13, -37	-38, -38, -76
R2	15, 32, 47	29, 43, 72	15, 26, 41	1, 1, 2	-13, -24,-37
R3	12, 18, 30	26, 29, 55	40, 40, 80	26, 15, 41	12, -10,2
R4	1,4, 5	15, 15, 30	29,26, 55	43,29, 72	29, 4, 33
R5	-10,-10,-20	4,1, 5	18,12, 30	32, 15, 47	46, 18, 64

Game 4

	C1	C2	C3	C4	C5
R1	14, 42, 56	0, 25, 25	-14, 8,-6	-28, -9, -37	-42, -34, -76
R2	11, 28, 39	25, 39, 64	11, 22, 33	-3, 5, 2	-17,-20,-37
R3	8, 14, 22	22, 25, 47	36, 36, 72	22, 19, 41	8, -6, 2
R4	5, 0, 5	19, 11, 30	33, 22, 55	47, 33, 80	33, 8, 41
R5	-6,-14,-20	8, -3, 5	22, 8, 30	36, 19, 55	50, 22, 72

Game 5

	C1	C2	C3	C4	C5
R1	10, 38, 48	-4, 21, 17	-18, 4,-14	-32,-13,-45	-46,-30, -76
R2	7, 24, 31	21, 35, 56	7, 18, 25	-7, 1,-6	-21,-16,-37
R3	4, 10, 14	18, 21, 39	32, 32, 64	18, 15, 33	4, -2, 2
R4	1,-4,-3	15, 7, 22	29, 18, 47	43, 29, 72	29, 12, 41
R5	-2,-18,-20	12,-7, 5	26,4, 30	40, 15, 55	54, 26, 80

Appendix D

Additional Regressions

Table D1: Regression Analysis of Messages

Dependent Variable	Message (Remapped)		Truth-telling	
Model Type	Ordered Probit		Probit (Marginal Effects)	
	Model 1	Model 2	Model 3	Model 4
Game (Remapped)	0.334*** (0.029)	0.309*** (0.045)	0.188*** (0.011)	0.163*** (0.016)
Strangers	0.304*** (0.065)	0.345* (0.189)	-0.134*** (0.044)	-0.246*** (0.071)
High State Losses	-0.178*** (0.080)	-0.711*** (0.250)	0.086 -0.054	0.206*** (0.079)
Rounds 10 - 18	0.145*** (0.044)	0.260** (0.112)	-0.017 -0.02	-0.151*** (0.043)
Game * Strangers		-0.015 0.049		0.037* (0.020)
Game * High State Losses		0.184*** (0.065)		-0.045** (0.023)
Game * Rounds 10 - 18		-0.040 (0.033)		0.046*** (0.013)
Log-likelihood	-4600.62	-4581.62	-2010.17	-1993.43
# Observations	3,564	3,564	3,564	3,564

Note: Standard errors, reported in parentheses, are corrected for clustering at the session level (strangers matching) or the group level (partners matching). Three (***), two (**), and one (*) stars indicate statistical significance at the 1%, 5%, and 10% levels respectively.

The regressions reported in Table D1 show that the relationship between the game being played and the message sent is statistically significant consistent with our discussion in Section 5.2.a. The dataset used for these regressions includes all observations from Baseline-Centralization, Strangers-Centralization, and HSL-Centralization. An observation is the message sent by a single division manager in a single round. The dependent variable for Models 1 and 2 is the message, remapped as described above for Table 7 (All observations are from a D1's point of view. Game 1 for a D2 is mapped to Game 5 for a D1, Game 2 for a D2 is mapped to Game 4 for a D1, etc). Remapping the games allows us to pool data for the two divisions into a single regression. Since messages have a natural order and are categorical, we use an ordered probit model in Models 1 and 2. For Models 3 and 4, the dependent variable is whether the division manager told the truth. This is a binary variable so we use a probit model, reporting marginal effects. In all regressions, the independent variables include the game (remapped for D2s), dummies for the Strangers and High State Losses treatments, and a dummy for Rounds 10 – 18. Models 2 and 4 add interactions between the game and the other three variables (dummies for Strangers, High State Losses, and Rounds 10 – 18). Observations are not independent, so the standard errors are corrected for clustering at the session level (strangers matching) or the group level (partners matching). This yields a total of 59 clusters.

Model 1 indicates that the relationship between the game being played and the message sent is strongly significant. This is the most important takeaway from the regressions, indicating that the messages convey information about the state of the world and supporting our description of the data in Section 5.2.a. Looking at Model 2, the relationship between games and messages is stable over time as indicated by the small and insignificant estimate for the interaction term between the game and Rounds 10 – 18. It is also worth noting that the relationship between the game being played and the message sent is significantly stronger in the High State Losses treatment.

Model 3 shows that truth-telling is significantly more likely for games closer to a division's preferred equilibrium. This makes sense as the divisions have less to gain by lying in these cases. The small difference in the frequency of truth-telling between the Baseline and Strangers treatments is significant, but the difference between the Baseline and High State Losses treatments is not. Looking at Model 4, the relationship between the likelihood of telling the truth and the game being played becomes more sensitive with experience. Subjects become more likely to lie when it is to their benefit and less likely to lie when gains from lying are small. The probability of telling the truth is significantly more sensitive to the game being played in the Strangers treatment than the Baseline treatment and significantly less sensitive in the High State Losses treatment than the Baseline Treatment. The latter result presumably follows from the change in payoffs which makes it less valuable to deceive the CM.

Table D2: Regression Analysis of Information Usage

	Model 1	Model 2	Model 3	Model 4
Dependent Variable	Row/Column	Row/Column	Row/Column	Row/Column
Game	0.207*** (0.030)			
D1 Message		0.406*** (0.027)	0.371*** (0.055)	0.323*** (0.034)
D2 Message		0.368*** (0.039)	0.400*** (0.052)	0.340*** (0.044)
Strangers			0.193 (0.338)	0.062 (0.107)
Strangers * D1 Message			0.022 (0.063)	
Strangers * D2 Message			-0.071 (0.081)	
High State Losses			-0.178 (0.430)	0.099 (0.096)
High State Losses * D1 Message			0.068 (0.087)	
High State Losses * D2 Message			-0.002 (0.092)	
Partners * Lagged Truth (D1)				-0.364 (0.256)
Partners * Lagged Truth (D2)				-0.541*** (0.186)
Strangers * Lagged Truth (D1)				-0.379 (0.326)
Strangers * Lagged Truth (D2)				0.094 (0.143)
Partners * D1 Message * Lagged Truth (D1)				0.138** (0.066)
Partners * D2 Message * Lagged Truth (D2)				0.127** (0.065)
Strangers * D1 Message * Lagged Truth (D1)				0.096 (0.061)
Strangers * D2 Message * Lagged Truth (D2)				-0.032 (0.056)
Log-likelihood	-2600.34	-2371.11	-2366.71	-2220.13
# Observations	1,748	1,748	1,748	1,658

Note: Standard errors, reported in parentheses, are corrected for clustering at the session level (strangers matching) or the group level (partners matching). Three (***), two (**), and one (*) stars indicate statistical significance at the 1%, 5%, and 10% levels respectively. Only observations where divisions' product types are coordinated ($T_1 = T_2$) are included in the regression.

This analysis reported in Table D2 confirms that information is transmitted from divisions to the CMs due to CMs responding significantly to the divisions' messages, with no significant difference between the three treatments (Baseline-Centralization, Strangers-Centralization, and HSL-Centralization). The dataset only includes observations where the CM chose the same product type for both divisions, imposing coordination ($T_1 = T_2$). This eliminates 34 out of 1782 observations, but allows us to summarize a CM's choice as a single number. The dependent variable is the row/column (1, 2, 3, 4, or 5) chosen by the CM. The CM's choices are categorical and ordered in nature, so we use ordered probit models. Standard errors are corrected for clustering at the session level (strangers matching) or the group level (partners matching).

If information is being transmitted from the divisions to the CM, the CMs' choices should be correlated with the game being played. In Model 1, the sole independent variable (other than round dummies) is the game being played. The estimate is positive and significant, indicating that successful information transmission occurs. Model 2 looks directly at the effect of the divisions' messages, with the messages of D1 and D2 added as separate variables. Both estimates are positive and significant, consistent with our observation from Table 3 that CMs respond to the divisions' messages. The two estimates are virtually identical (and not significantly different). There is no particular reason for CMs to be systematically more responsive to one division than the other, and indeed they aren't. Model 3 adds a dummies for the Strangers and High State Losses treatments as well as interactions between these dummies and the messages from the two divisions. None of these added variables are significant. There is no significant difference in responses to messages across the three treatments.

Model 4 looks at a different aspect of how CMs respond to divisions' messages. Recall that the feedback allows CMs to know, *ex post*, whether D1 or D2 lied. Model 4 modifies Model 2 by adding information about whether D1 or D2 told the truth in the previous round. In the Strangers treatment the CM knows that a lagged observation of truth-telling came from a *different* subject than the current division manager, but in either treatment with partners matching (Baseline and High State Losses) the CM knows that it came from the *same* division manager who sent the current message. We therefore expect the response to previous truth-telling to differ across the two types of matching. For each type of matching and each division manager (D1 and D2), we add a dummy for whether the CM observed truth-telling by that division manager in the previous round plus an interaction between that dummy and the division manager's current message (data from Round 1 are dropped since there is no previous round). This adds eight variables to the regression, two for each cell of matching type and division manager. We also add dummies for the Strangers and High State Losses treatments to avoid biasing the estimates for the terms relating to past truthfulness. For the Strangers treatment, none of the four variables related to past truthfulness are significant, but for the partners matching treatments, three of the four variables are statistically significant. For both division managers, the term for lagged truthfulness is negative and the interaction with the division manager's current message is positive. With partners matching CMs are more sensitive to a message from a division manager who they know told the truth in the previous round.

Table D3: Regression Analysis, Effects of Manager Suggestions on Equilibrium Type

Dependent Variable	Safe Equilibrium	Efficient Equilibrium
Suggest Safe Equilibrium	0.194** (0.082)	-0.059 (0.043)
Suggest Efficient Equilibrium (Game ≠ 3)	-0.356*** (0.063)	0.562*** (0.102)
Suggest Other Equilibrium	-0.270*** (0.076)	-0.119 (0.083)
Don't Suggest Equilibrium	-0.198** (0.079)	0.014 (0.068)
Lagged Safe Equilibrium	0.445*** (0.059)	-0.029 (0.035)
Lagged Efficient Equilibrium	-0.080* (0.042)	0.139** (0.058)
Lagged Other Equilibrium	-0.199* (0.106)	0.028 (0.131)
Log-likelihood	-406.96	-362.42
# Observations	918	918

Note: Data from Round 1 are dropped to allow use of lagged variables. Standard errors, reported in parentheses, are corrected for clustering at the group level. Three (***), two (**), and one (*) stars indicate statistical significance at the 1%, 5%, and 10% levels respectively.

The regressions reported in Table D3 show that the primary effect of realized suggestions of the safe equilibrium is to shift play away from the efficient equilibrium towards the safe equilibrium. These two regressions use the same specification and dataset as the regression in Table 5. The only change is the dependent variable; this is a dummy for the safe equilibrium in the first regression and a dummy for the efficient equilibrium in the second regression. For the first regression, the results are pretty much what would be expected. Suggesting the safe equilibrium makes the safe equilibrium more likely while suggesting anything else makes it less likely. The second regression is slightly less expected. A realized suggestion of the efficient equilibrium makes this outcome more likely, as expected, but other suggestions have no significant effect.

Appendix E

Table E1: Detailed Description of Coding

1) Make a suggestion about what row/column should be chosen. (Coder always recorded the specific suggestion that was made.)
    a. Suggest safe equilibrium
    b. Suggest efficient equilibrium
2) Agree to proposal about what row/column should be chosen.
3) Discussion about what row/column should be chosen.
    a. Discuss need for coordination (pick same row & column). This requires more than making a suggestion that involves coordination. The message needs to indicate that the two players should be choosing the same thing (e.g. "We'll do better if we make the same choices" is coded. "Let's choose row 4 and column 4" is not coded.)
    b. Discuss fairness. This category includes any message that discusses the distribution of pay over the three players
4) Discuss Efficiency: This includes discussion of maximizing total pay as well as explaining how and why rotation between players works.
5) Questions About Rules of the Experiment: This includes questions about either the rules of the experiment (e.g. "Do I choose a row or does [the central manager] choose for me?") or the game (e.g. "Is the third number my payoff?").
6) Questions About How to Play: This was for conceptual questions rather than the frequent generic request that somebody suggest a row and column.
7) Explanation: This included explanations about the rules of the experiment or game, as well as explanations of a suggested way of playing the game.
8) (Chat-Centralization Only) CM Asks What Game is Being Played.
9) (Chat-Centralization Only) Divisions Report What Game is Being Played.
    a. Truthfully Reveal Game
    b. Lie About Game
    c. Conflict: This is used for cases where there was "fact-checking". (e.g. "D1: It is Game 3." "D2: No, it is Game 2.") This category is different from the "Contradict" category reported in Table 6, which is a combination of 9a and 9b. This category is the basis for the discussion of fact-checking in the text.