

Collaborating at the Tower of Babel:

The Meaning of Cooperation and the Foundations of Long-Term Exchange

Hagay C. Volvovsky*

Dec. 2023

Abstract

Cooperation requires not only that exchange partners take cooperative actions but also that they arrive at a mutual understanding of what cooperation means – i.e., the specific actions that constitute cooperation and defection. Absent such shared understandings, even cooperative actions may be interpreted as defection, leading to the termination of the exchange. Past work has found that shared understandings of cooperation are present in long-term exchange relationships but has left open a fundamental question: if relationship longevity breeds shared understandings of cooperation, how can new exchange relationships that lack such understandings survive into longevity? This paper proposes and validates a general process by which exchange partners can arrive at a shared understanding of what cooperation means in changing and complex environments, thereby making both cooperation and long-term exchange possible. At the heart of the analytic strategy is the differentiation between two contexts: one in which exchange partners can rely on established, pre-existing, commonly known frameworks and one in which a mutual understanding cannot be achieved without the joint development of a new relationship-specific framework. A set of online experiments validates the key theoretical implication. When established frameworks are unavailable, low-stakes joint problem solving helps forge mutual understandings that support cooperation and long-term exchange. By contrast, when relationships are not needed to develop mutual understanding, joint problem solving (and long-term exchange), may not be valuable for supporting cooperation.

* Coller School of Management | Tel Aviv University. Email: hagayv1@tauex.tau.ac.il. I want to thank Abdullah Almaatouq, Patrick Barry, Lisa Bernstein, Rodrigo Canales, Roberto Fernandez, Carolyn Fu, Robert Gibbons, Amir Goldberg, Ko Kuwabara, Ray Reagans, J. Michael Wahlen, JoAnne Yates, and Ezra Zuckerman-Sivan.

“[T]he Lord said...’[T]hey have all one language... nothing that they propose to do will now be impossible for them. [L]et us... confuse their language, so that they may not understand one another’s speech.’ [S]o... they left off building the city. Therefore its name was called Babel, because there the Lord confused the language of all the earth.”

Genesis 11: 6-9

“And they forgot at once each other's speech, and they could no more speak nor understand the same language. And when the builder would receive from the hands of the carrier the wrong bricks, or something else he was not ordered to bring, the builder in his anger would fling it away, and cast it upon someone and kill him.”

Sefer ha-Yashar, Noah 14

1. Introduction

We have long known that long-term exchange relationships promote cooperation by reducing the risk of defection. Long-term exchange extends the shadow of the future, embeds exchange in supporting social relations, and generates trust (Axelrod and Hamilton, 1981; Granovetter, 1985; Gulati, 1995). Recent work on cooperation highlights an additional, oft-overlooked mechanism through which long-term exchange facilitates cooperation – helping exchange partners coordinate on what each ought to do to meet unexpected challenges and opportunities (Uzzi, 1997: 45–6; Gulati, Wohlgezogen, and Zhelyazkov, 2012). Such coordination is crucial, as exchange partners often have conflicting interests. As a result, exchange partners that fail to coordinate on what each should do in response to an unexpected event are prone to mistaking intendedly cooperative actions as acts of defection, leading to retaliation and dissolution of the relationship (Gibbons and Henderson, 2012). Cooperation, in other words, requires not only that

exchange partners take cooperative actions but also that they coordinate on what cooperation means – i.e., the specific actions that constitute cooperation and defection.

But if coordination on the meaning of cooperation is necessary for the survival of exchange relationships, and if long-term exchange is how coordination is achieved, we are left with a chicken-and-egg problem. Existing long-term exchange relationships are likely to endure, in part, because they foster coordination on the meaning of cooperation. New exchange relationships are less likely to survive to longevity precisely because they lack such coordination. How, then, can new long-term relationships be formed? What advice can we offer new exchange partners seeking to ensure their relationship does not collapse due to a coordination failure before longevity can be achieved?

Take, for example, Uzzi's (1997) account of the women's fashion industry in New York City. Uzzi found that designers preferred working with long-term manufacturing partners, in part because long-term ties allowed them to coordinate around the tacit and complex notions of "style" and acceptable costs. Long-term manufacturing partners could recognize unanticipated problems as they emerged and make adjustments that lived up to designers' expectations of how a dress should "fall," "run," or "catch light." And they could do so at a cost the designer would understand and accept. On the other hand, new manufacturing partners stuck to schematics and delivered unsatisfactory dresses and pricing data full of "funny numbers" that the designer did not understand or consent to (*ibid*, 45-7). Though left unsaid, it is likely such unsatisfactory results and pricing disputes will lead to the dissolution of many new designer-manufacturer relationships, preventing them from maturing into long-term partnerships.

To be sure, coordination on the meaning of cooperation is not always challenging. We can all agree that violating an explicit agreement, legally binding or not, is a form of defection or that an Uber driver must drive us safely to our destination. Coordination with an Uber driver is easy, not because we can agree in advance on the specific actions the driver must take – driving is far too unpredictable for that – but because we can rely on an established and widely shared tacit framework that coordinates our interpretation of an Uber driver’s actions as cooperative (or not). An established “coordination framework,” for short.

There are many circumstances, however, in which established coordination frameworks are largely unavailable. Idiosyncratic and incomplete contracts are the prime example: exchanges that are too unique to allow for the development of an established framework and too complex for exchange partners to agree in advance on what each ought to do in any likely eventuality. Diners in a restaurant cannot pre-specify precisely how they want their dishes to taste; partners in a strategic alliance cannot completely pre-specify what each ought to do for the alliance to succeed; and, as we have seen, designers and manufacturers cannot fully pre-specify what actions the manufacturer ought to take, and the designer ought to pay for. This leads to our question: How can exchange partners successfully coordinate (and thus be able to cooperate) in circumstances such as these?

Past work on cooperation has largely bracketed this question, assuming (often implicitly) that exchange partners coordinate successfully on the actions that constitute cooperation and defection (e.g., Axelrod and Hamilton, 1981; Kollock, 1994; Simpson and Willer, 2015). At the same time, work on long-term exchange relationships has shown that long-term exchange facilitates coordination. But as Uzzi’s study suggests, causation may run both ways: long-term

exchange may facilitate coordination, and successful coordination may facilitate long-term exchange. More importantly, the finding that (surviving) long-term exchange partners successfully coordinate offers little guidance to new exchange partners. Even if long-term exchange facilitates coordination, how can new partners ensure their exchange does not end prematurely due to a coordination failure? Below, I develop a first cut at addressing this question. The analysis yields fresh insights into what is necessary and sufficient for successful cooperation and expands our understanding of why and when long-term exchange relationships support cooperation.

At the heart of the analytic strategy is the differentiation between two contexts: one in which exchange partners can rely on an established coordination framework and, therefore, need not know each other to coordinate (and thus cooperate) successfully, and one in which coordination (and therefore cooperation) cannot be achieved without the joint development of a new, relationship-specific, mutually shared coordination framework – or “joint framework” – that helps exchange partners coordinate in the specific set of circumstances in which they find themselves. I proceed to identify the mechanism by which (surviving) long-term exchange relationships build joint frameworks, thereby making cooperation possible. I then show how new exchange partners can mimic this mechanism in the early stages of exchange relationships, increasing the odds that cooperation will succeed and the relationship will survive into longevity.

To anticipate the argument developed below, note that the content of coordination frameworks is inherently tacit. It can only partially be communicated in words and often requires prolonged interaction and shared experiences to be fully transmitted (Hansen, 1999; Reagans and McEvily, 2003). As a result, actors struggle to communicate them to strangers with whom they

have no shared history. Conversely, they find it much easier to communicate their coordination frameworks to long-term partners with whom they share a deep well of experiences (Uzzi, 1997; Hansen, 1999; Reagans and McEvily, 2003).

The most significant shared experience for the creation of joint frameworks is that of joint problem solving. As unanticipated circumstances inevitably arise throughout a long-term relationship, actors negotiate their implications, engage in joint problem solving, and, if successful, develop a shared language to discuss the problem and its solution and reach a shared understanding of what each ought to do (Uzzi, 1997). With time, these shared understandings accumulate to a reservoir from which shared interpretations of unexpected events can be drawn, and a joint framework is born. The difficulty lies in the early stages of this process. New partners lack shared experiences and joint frameworks and are thus prone to arrive at different interpretations of their obligations in the face of unexpected events. This creates the risk intendedly cooperative actions will be misinterpreted as acts of defection. As I explore below, new partners can mitigate this risk by opening exchanges with low-stakes joint problem solving that reduces the risk of defection (and misinterpretation).

I validate these propositions by developing a novel experimental strategy that can manipulate participants' coordination frameworks and assess the importance of joint coordination frameworks for successful cooperation. Throughout, I hold constant the risk of defection to focus on the effects of joint frameworks on cooperation, independent of the risk of defection. I apply this approach using a large web-based experiment in which participants are randomly assigned into one of two worlds: one in which coordination frameworks are established, and one in which they are constructed through low-stakes joint problem solving. Subsequently, participants are

engaged in a Principal-Agent game that exposes both to a risk of defection. The results demonstrate that when frameworks are established, low-stakes joint problem solving has no discernable effect on cooperation and participants show no discernable preference for repeated exchange with existing partners over exchange with new partners. But when participants cannot rely on established frameworks and have to jointly construct frameworks to successfully cooperate, participants cooperate more successfully with partners with whom they engaged in low-stakes joint problem solving in the past and prefer exchanging with those partners to new partners, thus opening the door to formation long-term exchange relationships. These results suggest that when established frameworks are unavailable, the creation of joint frameworks facilitates both cooperation and the formation of long-term exchange relationships.

The remainder of this paper is organized as follows. In section 2, I identify the theoretical challenge by highlighting coordination's critical yet overlooked role in successful exchange. I proceed to develop a theoretical framework explaining how long-term partners develop joint coordination frameworks to successfully coordinate when established frameworks are unavailable and how new partners can mimic this mechanism. I then advance a novel identification strategy for testing the effects of joint frameworks on cooperation. In section 3, I devise an experimental design to execute the identification strategy. In section 4, I discuss the results. In section 5, I discuss their implications.

2. Theory

The goal of this paper is to explain how exchange partners coordinate on the meaning of cooperation. Coordination's role as a necessary condition for successful exchange is an underappreciated thread in the literature on cooperation. Kreps (1990), for example,

conceptualizes organizational culture as a “focal point” (Schelling, 1960) that allows organizational members to coordinate on a shared understanding of what each ought to do. Uzzi (1997) describes the importance garment designers place on reaching shared interpretations of difficult-to-communicate “styles” and cost data with their manufacturers. Bernstein (2001) shows that shared quality evaluations are necessary for trade in the cotton industry. Gibbons and Henderson (2012) crystalize this insight; they show that actors’ ability to understand the content of each other’s promises – what they call *clarity* – underlies a wide range of exchanges and managerial practices, from the Toyota Production System to bonuses in Lincoln Electric and to Science-Driven drug discovery at Merck.

Yet, while useful, the only available account of how exchange partners can coordinate on the meaning of cooperation – the one offered by Gibbons and Henderson (2012) – suffers from significant limitations as an explanation for real-world coordination. Gibbons and Henderson posit that exchange partners coordinate on the meaning of cooperation by achieving “mutual knowledge” – that is, shared knowledge of all actions and payoffs available to them as acts of defection and cooperation (*ibid*, 1358). Mutual knowledge is attractive as an ideal first-best solution. If exchange partners could attain mutual knowledge, they could easily agree on which actions constitute cooperation. The core problem – that one person’s cooperation is another person’s defection – would thereby be eliminated.

However, this solution has two limitations as an explanation of how coordination is actually achieved. First, as Gibbons and Henderson recognize, gaining mutual knowledge of all available actions and payoffs is extremely difficult (*ibid*, 1358-60) – and is inherently impossible when it comes to incomplete contracts. By definition, parties to incomplete contracts cannot anticipate all

likely eventualities, let alone all the actions and payoffs available to them when unanticipated events occur (Hart and Moore, 1999). And since mutual knowledge often relates to unknowable events, it is impossible to measure and empirically validate. Suppose we wished to validate the claim that the attainment of mutual knowledge was responsible for a case of successful cooperation. To do so, we would have to demonstrate that the exchange partners achieved mutual knowledge of events neither they nor we could even imagine. As we cannot envision these events, the only way we could verify the exchange partners knew about them, is by observing the consequences of the exchange. Finally, this implies that accepting mutual knowledge as the explanation for how coordination is achieved results in the following tautological inference: if exchange partners coordinate successfully, they must have mutual knowledge; if they fail to coordinate, they must lack it. Mutual knowledge would, therefore, become the unfalsifiable *deus ex machina* of coordination (and thus cooperation) rather than a mechanism that can be identified empirically.

Second, we know that exchange partners can and do successfully cooperate (and hence coordinate) in a myriad of exchanges in which mutual knowledge is (as it must almost always be) unattainable. For example, long-term designer-manufacturer partners in Uzzi's (1997) study have demonstrated their ability to cooperate despite the significant risk of defection and the unattainability of mutual knowledge. The ever-changing nature of fashion styles and production line technology means that whatever mutual knowledge they arrive at any given time is of little use when unexpected events occur. The fact that these designer-manufacturer partners nevertheless found ways to successfully cooperate logically implies that they have found another solution to the coordination problem. Mutual knowledge, in other words, may be an ideal solution, but it is not a necessary one. An alternative, second-best solution must be at work in the

long-term designer-manufacturer relationships described by Uzzi, as well as other successful exchanges in which mutual knowledge is unavailable. But what is it?

I now develop a theory that answers this question in three steps. In the first step, I argue that joint coordination frameworks are a feasible and sufficient solution through which exchange partners coordinate on the meaning of cooperation when established frameworks are unavailable. In the second step, I identify a feasible way joint coordination frameworks may be achieved by demonstrating how they can emerge from long-term exchange via successful joint problem solving. I proceed to identify how new exchange partners can increase the odds joint problem solving will succeed when it is most perilous – at the beginning of a new exchange relationship – by lowering the stakes involved. This takes us to the third and final step, in which I devise an identification strategy to test joint frameworks' causal effect on cooperation. To that end, I analyze the conditions that affect the severity of the coordination problem in a given exchange. Since joint frameworks facilitate cooperation by mitigating coordination problems, their impact should be more pronounced the more severe the problem is. This means that we can identify the effect of joint frameworks on cooperation by comparing cooperation rates in exchanges that differ only in the severity of the coordination problem the parties face. Similarly, we can identify the effect of low-stakes joint problem solving on the creation of joint frameworks by comparing cooperation levels between partners who engage in it versus new partners in exchanges that differ only in the severity of their coordination problem. Together, these three steps resolve the twin problems of infeasibility and unidentifiability that afflict the mutual knowledge solution to the coordination problem. Further, they establish a novel explanation for how exchange partners coordinate on the meaning of cooperation and build a road map for empirically validating it.

2.1 Step One: Joint Coordination Frameworks

The key novel construct introduced in this paper is that of *coordination frameworks*, which I define as (inherently) tacit frameworks that guide actors' interpretations of what actions constitute cooperation and defection in the face of unexpected events (i.e., events for which exchange partners made no specific agreement). Unlike mutual knowledge, shared coordination frameworks are an attainable solution to the problem of coordination that can be identified empirically, at least in part.

While we are all familiar with established frameworks – such as those defining what an Uber driver and passenger must do during a ride[†] – joint frameworks are relationship-specific and thus largely hidden for those not party to the exchange. To get a glimpse into what joint coordination frameworks look like and how they can be empirically detected, consider Canales's (2014) study of three micro-credit financial institutions (MFIs) in Mexico. The loans offered by the MFIs are simple fixed-sum loans with pre-specified payment plans. Nevertheless, much like with the notion of “style,” their practical implementation requires complex coordination on issues such as when it is appropriate for a borrower to request a loan restructuring and when a loan officer should collect on a loan by seizing the borrower's assets or solicit private information and access to the borrower's business. The answers to these questions are highly subjective and situation-specific. Moreover, as in other environments we have reviewed, the setting is marked by

[†] Note that while the contract here is inherently incomplete – we cannot specify in advance the specific actions the driver ought to take to fulfil her end of the exchange – we can all agree that the driver must drive safely and quickly to the destination, and that the passenger must be cordial and keep the car clean. More importantly, for any particular set of actions by an Uber driver or passenger, different drivers and passengers can largely agree on whether those actions constituted cooperation.

unforeseeable difficulties (e.g., family health problems) and opportunities (e.g., a competitor going out of business), which routinely raise questions about what constitutes cooperation (vs. defection) and which render mutual knowledge unattainable. Successful cooperation thus requires coordination of borrowers' and loan officers' subjective evaluations of unanticipated events.

Seeking to explain how organizations manage the contradictory needs for flexibility and standardization in the face of such ambiguity using committees composed of both rule- and standard-oriented employees, Canales discovered that different loan officers within each of the MFIs exhibit one of two distinct “relational styles” in handling identical credit transactions (Canales, 2014: 4–5). “[L]etter of the law’ (LL) [loan] officers tend to adhere strictly to rules. They rely heavily on the standardized statistical tools that help automate decision making, strictly in line with organizational procedures” (*ibid*, 4). In contrast, the other relational style, the “[s]pirit of the law’ (SL) [loan] officers interpret policies expansively and flexibly. Rather than relying solely on rules, these officers establish embedded relationships with clients to gather fine-grained information and additional enforcement channels” (*ibid*, 4). These relationships often involved “[c]lose follow-up of business as well as personal activities, family issues, friendships, etc.” – all to get a better understanding of the borrower’s creditworthiness (*ibid*, 4–5).

The differences between the relational styles appear in stark relief whenever a borrower misses a payment. SL loan officers tend to be open to a client’s report that the missed payment resulted from an external or random shock and go out of their way to help these clients resume their capacity to service the loan. LL loan officers tend to interpret the same missed payments as

signals of low creditworthiness and dismiss borrowers' accounts of external shocks as "stories" crafted to take advantage of the MFI (*ibid*, 21-3). Presented with the same actions by borrowers, SL and LL officers arrived at very different interpretations of whether this action is an act of cooperation that increases the long-term value of the relationship or an act of defection, and act accordingly. Differently put, the relational styles Canales identifies at the individual level represent not only loan officers' orientation towards playing a role but also entail a framework each loan officer uses to interpret what cooperation means — i.e., what I am generally terming a coordination framework.

To successfully coordinate with borrowers despite this interpretive ambiguity, loan officers invest heavily in building joint frameworks that complement their relational styles. To that end, loan officers dedicate a significant part of their work to educating borrowers "on what it means to have a credit relationship" (*ibid*, 16). Loan officers are successful in that, in part, because most borrowers in this setting have little experience with the financial sector (*ibid*, 8). Borrowers, in other words, are a financial *tabula rasa* that loan officers can socialize into their interpretation of the role of loan officer and borrower, thus transforming loan officers' relational style into a joint coordination framework.

Subsequent analysis by Canales and Greenberg (2016) indirectly reveals just how vital those joint frameworks are to successful cooperation by examining what happens when they are disrupted. Seeking to explain how organizations can mitigate the loss of repeated exchange due to personnel turnover by maintaining a consistent role orientation, Canales and Greenberg find that when MFIs rotate loan officers, default rates increase whenever a loan officer exhibiting one relational style (and its associated coordination framework) replaces a loan officer with the other

relational style. This seems to occur because when loan officers and borrowers with different coordination frameworks are matched with each other, cooperative actions become more likely to be interpreted as acts of defection. Both parties are left unsure as to whether their cooperative actions will be interpreted as such. Coordination failure – and thus cooperation failure – ensues. This is true even when a flexible SL loan officer replaces a strict LL officer. As one borrower testified: “I really liked how she [an LL officer] worked. She came in, it was all business: Do you want to renew? How much? Any problems? Any referrals? And she was gone. That worked for me. But this new guy [an SL officer], man, he is so chatty. Always asking me about other stuff, and then telling me how to run my business. I just want my loan to run smoothly, as it did with [the previous loan officer]. So yeah, the change was not great for me” (*ibid*, 1209).

Finally, Canales (2014) and Canales and Greenberg (2016) have two limitations that help sharpen the agenda for this paper. First, both studies focus only on one side of the exchange: loan officers, not borrowers. Borrowers’ coordination frameworks are only observed indirectly through loan officers’ efforts to impart their framework on borrowers and the default rate spike that occurs when those efforts are disrupted. Second, while these studies give us confidence that joint coordination frameworks are a feasible, sufficient, and identifiable explanation for how exchange partners coordinate on the meaning of cooperation, they leave open the question of how exchange partners attain joint frameworks.

Fortunately, Canales and Greenberg (2016) also hint at a possible answer. By showing that joint coordination frameworks are often present in long-term exchange relationships and are difficult to transfer to new exchange partners, they suggest a link between long-term exchange and joint frameworks. At the same time, two additional limitations must be overcome to develop

and substantiate the relationship between long-term exchange and joint frameworks and advise new partners on building joint frameworks. First, the relationship between loan officers and borrowers is not symmetric. The borrowers in Canales (2014) have little choice but to adjust to loan officers' coordination frameworks. This raises the possibility that joint frameworks are only attainable in asymmetric long-term relationships.

Second, and more importantly, Canales and Greenberg (2016) only observe a correlation between (surviving) long-term exchange and joint frameworks indirectly and retrospectively through the consequences of its disruption. We cannot rely on such evidence to corroborate a causal link between long-term exchange and joint frameworks lest we risk falling into the same tautological reasoning described above with regard to relying on mutual knowledge to solve the problem of clarity: If the exchange is long-term, it must have succeeded and survived because exchange partners have a joint framework; if the exchange is short-term, it must have failed to survive because exchange partners lack a joint framework. Nor can we rely on the correlation between (surviving) long-term exchange and joint frameworks to help new exchange partners build joint frameworks. To escape this tautology and offer meaningful advice to new exchange partners, we need to unpack the link between long-term exchange and joint frameworks and explain how and when long-term exchange facilitates the creation of joint frameworks. I now do so in step two.

2.2 Step Two: The Origin of Joint Coordination Frameworks and the Value of Long-Term Exchange

Joint coordination frameworks are challenging to build for the very same reason they are necessary in the first place. Coordination frameworks come into play precisely in situations that

were neither agreed upon nor even conceived of by exchange partners. They cannot be reduced to clear-cut rules. Rather, they are general frameworks that help actors interpret actions as acts of cooperation versus defection in new and unexpected situations. Their content is, therefore, inherently tacit. As a result, the content of coordination frameworks cannot be effectively communicated in words and is thus extremely difficult to communicate to strangers. It can only be fully acquired through prolonged observation, interaction, and shared experience (Hansen, 1999; Reagans and McEvily, 2003).

Long-term exchange relationships can mitigate this problem. Studies of long-term exchange have shown that long-term exchange relationships often become embedded – that is, overlaid with social ties and characterized by close ongoing two-way interactions and shared experiences (Powell, 1990; Gulati, 1995; Uzzi, 1997; Elfenbein and Zenger, 2014). These shared experiences help exchange partners communicate the tacit content of their frameworks and converge on a shared one. Chief among them are shared experiences of joint problem solving. New and unexpected problems inevitably arise throughout the life of a long-term relationship. Exchange partners negotiate their implications and engage in joint problem solving to arrive at a shared interpretation of the situation and of what each of them ought to do (Uzzi, 1997; Hansen, 1999). If they fail to arrive at a shared interpretation, the relationship may end (Gibbons and Henderson, 2012). If they succeed, not only does the relationship survive, but as time passes and more problems are jointly solved, a reservoir of shared interpretations accumulates. This reservoir then provides exchange partners with ample sources of analogy from which shared interpretations of new situations can be drawn. A joint framework is born.

Recall, for example, Uzzi's (1997) finding that a key benefit of long-term exchange between designers and manufacturers is that long-term ties facilitate the communication of fine-grained tacit knowledge about “style” and “funny” pricing numbers (*ibid*, 45). This helps designers and manufacturers develop shared interpretative frameworks that allow them to construct shared interpretations of how dresses should look and the manufacturing adjustments the designer is willing to pay for to accomplish this look.

Uzzi's work falls into a broader category of scholarship, showing that much of the value of long-term exchange comes not from reducing the risk of defection but from smoothing interactions and allowing actors to focus on joint value creation (e.g., on the value of smoothing interactions in strategic alliances see Gulati, Wohlgezogen, and Zhelyazkov, 2012). Long-term relationships smooth interactions precisely by helping exchange partners coordinate what each should do to realize new opportunities and meet new challenges. This, in turn, increases the odds of successful cooperation, even when the risk of defection is not salient. A diner is likely to return to a restaurant she found to her liking, even if she has no reason to suspect chefs in other restaurants of shirking.

Such research on long-term relationships shows that they can foster the development of joint frameworks, but it gives little guidance to new exchange partners seeking to build long-term exchange relationships. How can a designer and manufacturer working together for the first time increase the likelihood of producing a satisfactory dress at an acceptable (and understandable) cost? Merely advising new partners to engage in long-term exchange is not enough. Shared experiences, after all, can also be negative. Joint problem solving can fail to produce an agreed-upon interpretation. Azoulay, Reppenning, and Zuckerman (2010) show that when exchange

partners have different understandings of cooperation and defection, the resulting relationships are “nasty, brutish, and short” (*ibid*, 499).

To see how new exchange partners can structure exchanges to increase the likelihood that joint problem solving will lead to joint frameworks and sustained cooperation, consider Bernstein's (2015) study of Midwestern original equipment manufacturers (OEMs). Bernstein finds that OEMs do not simply hand their suppliers a textbook detailing their expectations and proceed to dive right into an exchange relationship. Instead, they begin new exchange relationships with small, low-stakes transactions that significantly limit the incentive to defect. The reduction of defection incentives reduces the likelihood of conflict by increasing the likelihood that partners will cooperate and that any disagreements on the correct course of action in an unexpected situation will be interpreted as genuine rather than self-interested. It creates a space in which joint problem solving is likely to succeed, and joint frameworks can be formed (DiBenigno, 2020).

Bernstein finds that OEMs use these low-stakes exchanges to provide suppliers with on-the-job training, interactive feedback, and extensive joint problem solving. With time, as the OEM becomes satisfied that the supplier can meet the OEM's expectations, the OEM increases the scope and stakes of the transactions and engages the supplier in a full-blown exchange relationship (*ibid*, 578-82). While these practices are entirely consistent with the development of mutual knowledge (vs. joint frameworks) as a solution to the coordination problem – indeed, Bernstein often attributes their benefit to communicating information about OEMs' expectations – it is more likely that these practices build joint frameworks than shared knowledge. Much like fashion styles, technology and production processes at these OEMs and their suppliers are

constantly evolving. Mutual knowledge of all likely eventualities is simply unattainable; a joint coordination framework that helps exchange partners reach shared interpretations of unexpected situations is.

Finally, as illuminating as these examples are, they provide only suggestive evidence of a causal effect of joint coordination frameworks on cooperation. The critical issue is that the joint frameworks featured in these examples were developed endogenously between self-selected long-term exchange partners. Consequentially, two causal identification problems loom large. First, the observed relationship between joint frameworks and cooperation might be driven by partner selection. Perhaps exchange partners selected each other based on unobserved features and attributes that make cooperation more likely and which may be correlated with joint frameworks. If that is the case, the observed relationship between joint frameworks and cooperation is spurious. Second, extant research on long-term exchange relationships has identified a myriad of mechanisms by which long-term exchange can support cooperation. Consequentially, we cannot be sure whether cooperation in the examples above is driven by joint frameworks and not other correlates (or consequences) of long-term exchange.

To demonstrate that joint frameworks facilitate cooperation, coordination frameworks should be identified as causal even if they are not produced endogenously and are disentangled from other features of long-term exchange. This is the task of the third and final step of the proposed framework.

2.3 Step Three: Overcoming the Problem of Identifiability

Ideally, to test the causal link between joint frameworks and cooperation, we would want to randomly assign pairs of actors into exchange relationships and then compare cooperation levels within the same relationship with and without joint frameworks. Unfortunately, it seems impossible to experimentally insert and remove joint frameworks from exchange relationships. We can, however, get around this hurdle by identifying the conditions that induce exchange partners to invest in joint frameworks. We can then employ an intent-to-treat design to construct functionally identical exchange relationships that differ only in the extent to which exchange partners are likely to build joint frameworks. To that end, we need to take a step back and consider the conditions under which coordination is easy. Under such conditions, joint frameworks are of little use, and exchange partners have little incentive to invest in developing them. We can then exploit the variation between exchanges with more or more simple coordination problems to create identical exchange relationships that differ only in the extent to which exchange partners develop joint coordination frameworks.

It is thus helpful to array exchange situations in terms of the difficulty of the coordination problem they pose. Two factors determine the severity of the coordination problem: (1) the degree to which exchange partners can rely on established coordination frameworks and (2) the extent to which exchange partners have developed a relationship-specific joint coordination framework.

At the most basic level, many exchanges can rely on an established and widely diffused coordination framework, rendering the coordination problem almost trivial. We all know what cooperation means in these situations and what each exchange partner should do in any likely eventuality. Broadly speaking, these kinds of exchanges tend to be instantaneous, generic, and

common. The more instantaneous the exchange, the less likely unanticipated events will occur. The more generic and common it is in a particular culture, the more likely broadly diffused frameworks will emerge.

Consider, for example, purchasing a box of cereal from a supermarket. The exchange is instantaneous and largely generic: few, if any, actor-specific adjustments are required to make the exchange work. It is also ubiquitous in developed countries and has been widely represented in movies, books, and other forms of cultural expression. As a result, the vast majority of consumers in developed countries agree on what they must do to execute the exchange: pick up the cereal box, stand in line for the cashier, and pay the listed price. But it is only trivial within a culture that features established frameworks exchange partners can rely on. A visitor from Mars or a buyer from a different culture that does not have supermarkets might have a completely different coordination framework. He might, for example, think the goods in the supermarket are free, walk to the front of the line, or haggle with the cashier (Zuckerman Sivan, 2017).

Conversely, the more prolonged and idiosyncratic the exchange, the less exchange partners can rely on established frameworks to coordinate their understanding of what actions constitute cooperation and defection. Such is the case with many economically and organizationally significant exchanges, from the designer-manufacturer exchanges in Uzzi (1997) to strategic alliances in Biotech and executive hires. All are prolonged, actor or organization-specific, and thus idiosyncratic exchanges that can, at best, only partially draw on established frameworks from the general, industry, or professional culture. Instead, they must rely on joint frameworks or (oft unattainable) mutual knowledge developed by the exchange partners themselves (e.g., strategic alliances). Simply put, while many designers design dresses, each has a different style.

Consequently, the specific problems that may occur between a given designer and its manufacturer are (at least in part) unique to the designer. No widely diffused established framework can guide designers and manufacturers in interpreting them.

This brings us to the second factor determining the severity of the coordination problem: whether exchange partners share a relationship-specific coordination framework. When established frameworks are unavailable, coordination becomes much more difficult. Exchange partners can mitigate this difficulty by developing a joint coordination framework. Finally, coordination is most difficult when exchange partners can neither rely on an established or joint framework.

----- Table 1: Severity of the Coordination Problem -----

The different degrees of difficulty of the coordination problem allow us to experimentally manipulate the creation of joint frameworks in otherwise identical exchange relationships by manipulating the severity of the coordination problem. As Bernstein (2015) demonstrates, developing joint frameworks is a lengthy and expensive process. As a result, when coordination is easy because established frameworks are available, even long-term exchange partners are unlikely to invest in creating joint frameworks. However, when coordination is difficult because established frameworks are not available, exchange partners have a strong incentive to invest in developing joint frameworks to facilitate exchange. This means that exchange partners who cannot rely on established frameworks will be far more likely to develop a joint framework.

Having established the conditions under which exchange partners are likely to invest in creating a joint framework, we need to give them an opportunity to do so. Recall that (1) the

primary mechanism by which (surviving) long-term exchange facilitates joint frameworks is a shared history of joint problem solving; (2) that a key hurdle for the development of joint frameworks (and long-term exchange) between partners with conflicting interests is that joint problem solving may fail to produce an agreed upon interpretation, leading to the termination of the relationship before joint frameworks can form; and (3) that the OEMs in Bernstein (2015) overcome this hurdle by beginning relationships with low-stakes exchanges. Taken together, this suggests that we can facilitate the creation of joint frameworks between new exchange partners – like those we experimentally create – by initiating relationships with low-stakes exchanges (and, thus, low-stakes joint problem solving).

The last step of our identification strategy is to disentangle the effect of joint frameworks on cooperation from potential confounds. We can disentangle the effect of joint frameworks on cooperation from other correlates and outcomes of low-stakes joint problem solving, such as trust or positive affect (Lawler, Thye, and Yoon, 2009; Molm, Schaefer, and Collett, 2009), by comparing the effect of joint frameworks on cooperation across identical pairs of exchange partners who engaged in joint problem solving under conditions that do (or do not) induce joint frameworks. If joint frameworks facilitate cooperation, exchange relationships formed under conditions that induce joint frameworks should facilitate cooperation more (vs. exchange with new partners) than relationships formed under conditions that do not induce joint frameworks. Finally, we can disentangle the effects of joint frameworks from the risk of defection by holding the incentive structure constant across all exchanges.

Thus, I hypothesize[‡]:

Hypothesis I: *Holding the risk of defection constant, exchange partners who had engaged in low-stakes joint problem solving will successfully cooperate more often than new exchange partners.*

Hypothesis II: *Holding the risk of defection constant, the cooperation advantage enjoyed by exchange partners who had engaged in low-stakes joint problem solving will be greater when the exchange cannot rely on established frameworks.*

Hypothesis I is in line with the standard and well-researched insight that a history of prior exchange (and associated joint problem solving) facilitates cooperation. Hypothesis II breaks down this insight and posits a clear and novel mechanism, namely that the effect of prior exchange that involves joint problem solving on successful cooperation is driven, at least in part, by joint frameworks.

3. Data and Methods

3.1 Experimental Design

[‡] You can find the pre-registration of the hypothesis at: https://aspredicted.org/blind.php?x=4KM_K1B. Note that while the structure of the experiment and the hypothesis are as described in the pre-registration, there are two notable differences: First, following feedback from colleagues, I have changed the term “same norms” to “joint frameworks” to better reflect the construct of interest more accurately. Second, at the time of the pre-registration, I intended to run two related experiments. I only report on the first in this paper, corresponding to Hypothesis I in the pre-registration. The experiment corresponding to Hypothesis II was discontinued due to lack of power and will be re-run in the coming months.

To test whether joint coordination frameworks facilitate cooperation, I use an experiment that employs a cross-subject intent-to-treat design composed of two studies: (1) Study 1 manipulates the creation of joint coordination frameworks, and (2) Study 2 tests the effect of joint frameworks on cooperation. Both studies were developed using the Empirica platform (Almaatouq et al., 2021).

Specifically, Study 1 couples the insight that low-stakes joint problem solving facilitates the creation of joint frameworks with the insight that exchange partners are far more likely to develop joint coordination frameworks when established frameworks are unavailable. Participants in this study are randomly assigned to three-person groups and asked to play a variation of the standard coordination game commonly used to study the diffusion of social conventions and categorization systems (e.g., Weber and Camerer, 2003; Centola and Baronchelli, 2015; Burt, Reagans, and Volvovsky, 2021; Guilbeault, Baronchelli, and Centola, 2021). This task was selected for three reasons. First, it involves low-stakes joint problem solving. Second, coordination frameworks are a special case of a categorization system that categorizes actions as cooperation or defection. Third, using a variation on a standard task ensures the robustness of the experimental design and increases the legibility of the results.

Unlike the standard coordination game, which features a difficult coordination problem, Study 1 manipulates participants' frameworks by creating two parallel worlds: one featuring an easy coordination problem and one featuring a difficult one – that is, one with and one without established coordination frameworks. Participants in three-person groups that cannot rely on established frameworks have a much higher incentive to create joint frameworks and are thus much more likely to develop them. Finally, consistent with prior research, these joint

frameworks – participants’ solutions to the coordination problem – are likely to be group-specific and vary across different three-person groups (e.g., Weber and Camerer, 2003). As such, they are likely to facilitate coordination within groups but not across them. Participants in three-person groups that can rely on established frameworks, on the other hand, will have a far lower incentive to develop coordination frameworks and will thus be far less likely to create group-specific joint frameworks.

Study 2 builds on Study 1. In this study, participants from Study 1 are randomly assigned to Principal and Agent roles and play a variation on a standard cooperation game. Much like a standard cooperation game, both principal and agent have incentives to defect, and both benefit from successful cooperation. Unlike a standard cooperation game, successful cooperation in Study 2 requires successful coordination on the actions that constitute cooperation and defection. Conditions in Study 2 differ based on whether the Principal and Agent have engaged in low-stakes joint problem solving by playing together in Study 1, and whether established coordination frameworks are available. This results in a 2x2 design:

----- Table 2: Experimental Conditions -----

Throughout the two studies, participants are rewarded with points, which are converted to dollars upon completion.

3.1.1 Study 1: The Frameworks Game

Coordination frameworks are tacit and complex frameworks that take time and effort to build. As such, it is difficult to directly manipulate them experimentally or even detect their existence. Indeed, the only study that observed coordination frameworks, Canales (2014),

required deep ethnographic work, which is not practicable experimentally. That may be the reason for the paucity of systemic research in this area.

To address this problem, I leverage prior experimental work on the emergence of coordination conventions to create a novel experimental platform that indirectly shapes participants' frameworks by affecting how they make sense of the world. Specifically, Study 1 manipulates the creation of coordination conventions participants use to make sense of experimental stimuli. To that end, participants are randomly assigned to three-person groups, each of which is placed in one of two experimental conditions: (a) one in which participants can rely on an established framework (Established Frameworks: Available) and (b) one in which they cannot (Established Frameworks: Unavailable). Following experimental literature on the creation of coordination conventions, the game takes the form of an image-naming game (Weber and Camerer, 2003; Centola and Baronchelli, 2015; cf. Burt, Reagans, and Volvovsky, 2021).

In each round of the game, participants in three-person groups are shown a different image and asked to form a social convention on how the image will be named. Each participant may propose a name. The remaining participants may accept the name or reject it. Throughout, participants can communicate with each other using a three-person chat box (see Figure 1). The name that emerges constitutes a shared coordination convention that participants use to make sense of the image. Once participants agree on a name, they are rewarded and move on to the next round. Participants repeat this process for seven rounds, each featuring a unique image. The seven images used in each experiment run are randomly drawn from a set of twelve images.

----- Insert Figure 1 Here -----

Experimental conditions differ in the set of images from which images are drawn. The Established Frameworks: Available condition features standard images with established naming conventions such as an X sign, a star, a heart, etc. (see Figure 2). Established Frameworks: Unavailable condition features amorphously shaped symbols for whom there are no established naming conventions drawn from Wilkes-Gibbs and Clark (1992) (see Figure 3).

----- Insert Figure 2 and Figure 3 Here -----

By the end of Study 1, each group creates seven coordination conventions. Together, these shared coordination conventions constitute a coordination framework participants use to make joint sense of the experimental stimuli.[§]

3.1.2 Study 2: The Cooperation Game

Study 2 employs a variation on the standard Principal-Agent design. Participants are randomly assigned to a Principal or an Agent role and one of two experimental conditions: (1) Principal and Agent have engaged in low-stakes joint problem solving by playing together in the same three-person group in Study 1; (2) Principal and Agent are complete strangers.

To ensure that participants can tell whether they have engaged each other in joint problem solving, each participant is assigned a unique animal avatar, and each Study 1 group is assigned a

[§] I employ several measures to ensure data quality. Before beginning the game, participants are instructed about the game's structure and answer a series of comprehension checks. As an additional manipulation check, after completing seven rounds, participants repeat the process for a second time. As before, they play seven different rounds with the same seven images. This time, however, participants cannot communicate with each other and must each enter a name within 30 seconds. If all three names match, participants are rewarded and proceed to the next round. If the names do not match, participants repeat the same round with the same image.

unique color. Throughout the two studies, participants can see their color-avatar combo and that of the participants they are playing with. Thus, avatars serve as unique identifiers, and matching colors indicate that both participants played together in Study 1. Data collected from the exit survey suggests that participants could tell whether they played together in Study 1.

Each round of Study 2 is composed of the following stages: (1) the Principal decides whether to hire the Agent at the cost of 50 points or terminate the experiment (see Figure 4); (2) if the Principal hires the Agent, the Principal is shown a randomly selected image from the set of seven images used in the Study 1 (henceforth: Principal's Image) and asked to communicate the image to the Agent using a textbox (hereafter: the Principal's Message) (see Figure 5); (3) the Agent chooses whether to spend 30 points out of the 50 received from the Principal to read Principal's Message. If the Agent chooses to pay, the Principal's Message is revealed. If the Agent refuses (i.e., defects), the message remains hidden (see Figure 6); (4) the Agent is then presented with six different images randomly drawn from the images used in Study 1. One of the images is the Principal's Image. The Agent then selects an image (see Figure 7); (5) If the Agent picks the Principal's Image, the Principal receives 80 points. If the Agent chooses another image, the Principal receives nothing; (6) The Principal observes the outcome (reward/no reward) but not the Agent's choices (see Figure 8). As a result, the Principal *cannot directly observe* whether a failure – i.e., the Agent picking the wrong image – is caused by defection or an honest mistake stemming from differences in their coordination frameworks; (7) the Principal observes the outcome and chooses whether to hire the same Agent for another round or terminate the experiment (see Figure 9 for a game tree).

----- Insert Figures 4-9 Here -----

Note that having been paid 50 points by the Principal, the only incentive for the Agent to pay 30 points to read Principal's Message is that selecting the Principal's Image increases the odds of being hired for an additional round. If the Principal decides to nevertheless terminate the experiment, the Agent gains nothing for the 30 points spent on reading the Principal's Message. Thus, the Agent must select between cooperating by paying to read the Principal's Message, hoping that the Principal will reciprocate by hiring the Agent for another round, or defecting by refusing to pay for the message.

Conversely, the Principal gets rewarded (80 points for a net gain of 30) only if the Agent picks the Principal's Image. The Principal can thus decide to cooperate by hiring the Agent and risk losing 50 points or defect by terminating the experiment and depriving both participants of the gains of exchange. If the Principal chooses to cooperate and hires the Agent, and the Agent picks any image but the Principal's Image, the Principal loses the 50 points paid to the Agent and gains nothing. As a result, as far as the Principal is concerned, a failure to pick the Principal's Image constitutes defection by the Agent. Finally, note that the Principal cannot observe whether the Agent paid to read the Principal's Message and thus cannot distinguish a failure due to the Agent's refusal to pay for the message from a failure due to the Agent misunderstanding the message. Therefore, the actions and payoffs available for Principal and Agent as cooperation and defection are as follows:

----- Table 3: The Meaning of Cooperation and Defection -----

This structure of actions and payoffs is identical for all experimental conditions. As a result, the incentive for (and thus the risk of) intentional defection is held constant throughout the

experiment. And, since both principal and agent are equally dependent on each other for successful cooperation, there are no power differentials between the two roles.

To ensure data quality, before beginning the game, the Principal and Agent are instructed about each other's decisions and payoff structure, answer a series of comprehension checks, and play a training round. To prevent unraveling due to end-game effects, participants are told that the likelihood the Principal will be able to hire Agent for an additional round after one has been concluded is 70%. In practice, to maintain uniformity, a random draw with a 70% chance of success was used to determine a maximum of five rounds.

Finally, upon completion of Study 2, participants are directed to an exit survey collecting basic demographic data (age, gender, ethnicity, education) and the following open-ended questions: (1) how would you describe your strategy in the game? (2) Were the Principal and Agent in Game 3 [Study 2 in this paper] members of the same Game 1-2 [Study 1] group? How did that affect your strategy? (3) Do you feel the pay was fair? (4) We try to make this experiment as fun and engaging to participants as we can and would love any feedback you may have on how to improve it. Do you have any feedback for us, or would like to report any problems you encountered?

3.2 Participants

1,070 participants were recruited using Amazon Mechanical Turk (MTurk) from September 2021 to January 2022. Studies of the participant pool of MTurk have shown it to be a reliable source for social science research that is more representative of the US population than traditional lab studies (Berinsky, Huber, and Lenz, 2012; Paolacci and Chandler, 2014).

Participation was limited to US-based adults whose first language is English. A typical run of the experiment featured 6-12 participants and ranged from 15-25 minutes (average 21.6 minutes). Participants were paid \$4.5 on average, resulting in an average hourly pay of over \$12/hour. Participant demographics are displayed in Table 4.

----- Table 4: Participant Demographics -----

3.3 Outcome Variables

The hypotheses test the conditions under which actors are more likely to cooperate successfully. The natural outcome variable to measure successful cooperation is the number of successful rounds of cooperation. A successful round of cooperation is defined as a round in which both participants benefit. That is a round in which the Principal hires the Agent, and the Agent selects the Principal's Image. Since participants may play up to five rounds with each other, the number of successful rounds ranges from 0-5. Figure 10: Distribution of Successful Rounds displays the distribution of successful rounds played by the four experimental conditions.

----- Insert Figure 10 -----

3.4 Analysis

To test the hypotheses, I use an OLS regression to estimate the effects of established frameworks and a history of past exchanges involving low-stakes joint problem solving (consists of playing together in Study 1) on participants' decisions to engage in exchange in the Cooperation Game and their success in doing so. The following specification is used:

$$\begin{aligned}
Cooperation_i = & \beta_0 + \beta_1 \times Joint\ Problem\ Solving_i \\
& + \beta_2 \times Established\ Frameworks: Unavailable_i \\
& + \beta_3 \times Joint\ Problem\ Solving_i \times Established\ Frameworks: Unavailable_i
\end{aligned}$$

Where $Cooperation_i$ equals the number of successful rounds of cooperation by Principal-Agent pair i ; $Joint\ Problem\ Solving_i$ is a dummy variable equal to 1 if the Principal and Agent pair i have engaged each other in low-stakes joint problem solving in Study 1 and 0 otherwise; and $Established\ Frameworks: Unavailable_i$ is a dummy variable equal to 1 if the images used in both games were drawn from the Established Frameworks: Unavailable condition and 0 if they were drawn from the Established Frameworks: Available condition.

Thus, β_1 tests the first hypothesis and indicates the baseline effect of low-stakes joint problem solving on successful cooperation. β_2 indicates the baseline effect of being able to rely on established coordination frameworks. As such, it serves as a test of the manipulation in Study 1. A negative value for β_2 will indicate that the manipulation was successful – that is, coordination (and thus cooperation) was more difficult when established frameworks were unavailable. β_3 tests the second hypothesis. Finally, β_2 tests the second hypothesis. A positive value for β_3 will indicate that low-stakes joint problem solving is more likely to promote cooperation when it is more likely to lead to the creation of joint frameworks.

4. Results

4.1 Study 1: The Frameworks Game

Consistent with prior work, different three-person groups converged on different coordination conventions and created different coordination frameworks. This was especially so

in the Established Frameworks: Unavailable condition, which features images without established naming conventions. On average, different groups in this condition created 18.5 unique naming conventions for each image, compared to only 3.6 unique naming conventions per image in the Established Frameworks: Available condition.

The most common naming convention for each image in the Established Frameworks: Unavailable condition was on average shared by only 30% of the groups (if we remove two outlier images, that goes down to 21%), compared to 79% in the Established Frameworks: Available condition. To give but one example, the image depicted in Figure 11 appeared in 25 different games in the Established Frameworks: Unavailable condition. These 25 groups created 16 unique naming conventions, such as “Frankenstein,” “dancer,” and “penguin.” The most popular naming convention – “dancer” – was adopted by only 24% of teams.

----- Insert Figure 11 Here -----

The naming conventions, and thus coordination frameworks, created in the Established Frameworks: Unavailable condition were not only different but often mutually illegible. This can be seen both from the discordant naming conventions – e.g., the image in Figure 11 being named “dancer” and “Frankenstein” by different groups – as well as the fact that different groups used the same naming convention for different images. For example, the naming convention most commonly used for the image in Figure 11 – “dancer” – was also used by 22% and 21% of the groups for the images in Figure 12 and Figure 13.

----- Insert Figure 12 and Figure 13 Here -----

By contrast, the naming conventions, and thus coordination frameworks used in the Established Frameworks: Available condition, were mutually legible. Rather than creating new names, participants relied on existing, established names for the images. Even the image that generated the most unique names in this condition – the X sign – produced mutually legible names such as “x,” “wrong,” and “stop.”

4.2 Study 2: The Cooperation Game**

Figure 14 displays the average number of successful rounds by experimental condition. Error bars are set at 95% confidence intervals. Figure 14 reveals that Study 1 successfully manipulated the severity of the coordination problem. As expected, cooperation was far more likely when the coordination problem was simpler – that is, when participants could rely on established frameworks. Figure 14 also provides support for the two hypotheses. In line with Hypothesis I, the comparison between the orange and blue bars reveals that exchange partners who engaged in low-stakes joint problem solving are more likely to cooperate successfully than new partners. And in line with Hypothesis II, the figures reveal that low-stakes joint problem solving is more valuable (relative to new relationships) under conditions that induce the creation of joint frameworks. This can be seen from the fact that low-stakes joint problem solving facilitates cooperation in the left panel – i.e., when exchange partners develop joint frameworks because established frameworks are unavailable – but not in the right panel – i.e., when even exchange partners have no reason to invest in joint frameworks.

** Data and analysis code can be found at: https://researchbox.org/1849&PEER_REVIEW_passcode=FFCNOX

----- Insert Figure 14 Here -----

Before proceeding to the regression results, note that the experimental conditions were imbalanced due to a related experiment not reported here.^{††} To re-balance the conditions, 112^{‡‡} Principal-Agent pairs were selected at random from each of the four experimental conditions for the analysis presented below. Results are robust to using the complete data set.

^{††} The omitted experiment was focused on testing the effects of coordination frameworks on the spread of reputational information. It involved an observer that observed the Cooperation Game and was then asked whether she would like to replace the Principal and play with the Agent. Principal and Agent were not aware of the existence of the Observer and did not interact with her during the Cooperation Game. Unfortunately, technical issues with the Observer role combined with limited resources and the fact the analysis involved a triple interaction term prevented me from achieving sufficient power for this experiment.

^{‡‡} 112 was the smallest number of observations per condition.

Table 5 shows the OLS regression results. Model 1 tests the hypotheses. First, note that the results indicate that the manipulation in Study 1 was successful. Exchange partners cooperate less when the coordination problem is more complex – i.e. when exchanges cannot rely on established coordination frameworks. Principal-Agent pairs who could rely on established frameworks cooperate on almost two rounds more than those who could not ($\beta_2 = -1.98, P_{value} = 0.00$). This effect is statistically and economically significant, nearly doubling the rates of successful cooperation. Pairs who could not rely on established coordination frameworks cooperated on only 2.51 rounds. Pairs that could rely on established frameworks cooperated successfully on 4.48 rounds.

More importantly, the results of Model I support Hypothesis II but not Hypothesis I. In line with Hypothesis II, they indicate that the value of low-stakes joint problem solving is much greater when exchange partners are likely to invest in joint frameworks. At the same time, low-stakes joint problem solving, in and of itself has no statistically significant effect on cooperation. Its effect on cooperation is channeled entirely through the creation of a joint coordination framework. That is, only when exchanges are formed under conditions that foster the creation of joint frameworks – when established frameworks are unavailable – does a history of low-stakes joint problem solving support cooperation. The results indicate that partners with a history of low-stakes joint problem solving that induced them to create joint frameworks cooperated, on average, in 1.25 more than new partners in equivalent exchanges ($\beta_3 = 1.25, P_{value} = 0.00$). Given a base rate of only 1.76 successful cooperation rounds between new partners that cannot rely on established frameworks, a history of low-stakes joint problem solving improves cooperation by 71% when formed under conditions that encourage the development of a joint framework. Finally, by demonstrating that exchange relationships formed under conditions that

induce joint frameworks facilitate more cooperation than identical exchange relations that do not, these results provide evidence that joint frameworks facilitate cooperation.

Models 2-4 of Table 5 perform an exploratory analysis of possible mechanisms through which joint frameworks support cooperation. Model 2 examines the effect of joint frameworks on Principals' decisions to engage in exchange. To that end, the model replaces *Cooperation_i* with *Rounds Played_i* as a dependent variable. The results indicate that participants are aware of the severity of the coordination problem and are much more reluctant to enter exchange when coordination is more complex. Principals who could not rely on established coordination frameworks entered into 1.179 fewer exchanges than principals who could rely on them ($\gamma_2 = -1.179, P_{value} = 0.00$). Model 2 also shows that exchange partners in relationships formed under conditions that engender the development of joint frameworks were more willing (vs. new partners) to engage in exchange with each other ($\gamma_3 = 0.89, P_{value} = 0.012$). This indicates that principals were either aware of or learned to appreciate the value of joint frameworks for successful exchange. Finally, as in Model 1, Model 2 demonstrates that a history of low-stakes joint problem solving, in and on itself, has no discernable effect on participants' willingness to enter exchange with each other. It was only when exchange relationships were formed under conditions that encouraged joined frameworks that principals showed a clear preference towards hiring agents with whom they have engaged in low-stakes joint problem solving in the past, thus opening the door for the formation of a long-term exchange relationship.

Model 3 looks at the other side of the exchange and explores the effect of coordination frameworks on agents' (intentional) defection decisions. The more complex the coordination problem, the more difficult it is to distinguish deliberate defection from a coordination failure.

Thus, actors may take advantage of the interpretive ambiguity present when the coordination problem is more complex to defect more often. To explore this possibility, Model 3 replaces $Cooperation_i$ with $Effort Rate_i$ as a dependent variable. $Effort Rate_i$ is calculated by dividing the number of rounds the Agent paid to read the Principal's Message by the number of rounds played. It represents the likelihood that an Agent will cooperate in a given exchange. Using $Effort Rate_i$ as an outcome variable reduces the sample to 382 Principal-Agent pairs as 66 Principals chose to terminate the experiment at the beginning of the first round. The results show limited evidence that exchange partners are more likely to defect when the coordination problem is more complex. Agents engaged in exchanges that could not rely on established frameworks were less likely to cooperate ($\delta_2 = -0.068, P_{value} = 0.069$).

Finally, Model 4 attempts to disentangle the decision to engage in exchange from the exchange's success or failure. To that end, Model 4 replaces $Cooperation_i$ – a measure of the number of successful cooperation rounds – with $Success Rate_i$ as dependent variable. $Success Rate_i$ is calculated as the number of successful rounds played by pair i divided by the number of rounds played by pair i . It reflects the probability pair i successfully cooperates conditional on engaging in an exchange. The results reported in Model 4 are fully consistent with Model 1. Together, Models 2 & 4 indicate that joint frameworks facilitate cooperation both by increasing actors' willingness to enter into an exchange with each other and the likelihood the exchange will be successful.

----- Insert Table 5 -----

5. Discussion

To recall, the main objective of this paper is to explore the conditions under which exchange partners can successfully cooperate. I argued that cooperation requires not only that exchange partners overcome the temptation to defect and follow through on their promises but also that they coordinate on what those promises mean – i.e., the actions that constitute cooperation and defection. To date, the only explanation offered for how exchange partners coordinate is that coordination is achieved by building mutual knowledge of all available actions and payoffs (Gibbons and Henderson, 2012: 1358). I have shown that this explanation suffers from two fundamental limitations: it is often infeasible for exchange partners, and it is empirically unidentifiable. Yet the fact that exchange partners cooperate successfully without mutual knowledge indicates the existence of alternative solutions to the coordination problem.

To address these limitations and explain how real-world exchange partners cooperate, I developed and tested a novel alternative explanation in five steps: (1) drawing on Canales's and Uzzi's in-depth studies of successful exchange, I argued and shown that joint coordination frameworks are a feasible, sufficient, and identifiable solution to the problem of coordination employed by real-world exchange partners to solve the problem of coordination (Uzzi, 1997; Canales, 2014; Canales and Greenberg, 2016). More specifically, I have demonstrated that joint coordination frameworks help actors reach shared interpretations of what actions constitute cooperation and defection when unanticipated events inevitably occur; (2) I identified how exchange partners can develop joint coordination frameworks by engaging in (successful) joint problem solving in exchanges that cannot rely on established coordination frameworks. Of course, joint problem solving between exchange partners with conflicting interests may very well fail to produce a shared interpretation of cooperation, leading to the termination of the relationship – especially when it comes to new exchange partners that have no reservoir of

shared interpretations to draw on; (3) I therefore proceeded to identify steps new exchange partners can take to increase the likelihood joint problem solving succeeds and joint frameworks are created: beginning new relationships with low-stakes exchanges; (4) I then analyzed the conditions that affect the severity of the coordination problem in a given exchange and thus the importance of joint coordination frameworks; and (5) building on the above, I devised a novel identification strategy and experimental platform to test the causal effect of joint coordination frameworks on cooperation. The results indicate that when the coordination problem is complex – as in many economically and organizationally significant exchanges – joint coordination frameworks developed through low-stakes joint problem solving significantly increase the likelihood of successful cooperation and repeated exchange.

In doing so, my main contribution is clarifying the foundations of cooperation by ruling in a feasible and identifiable explanation for how exchange partners coordinate on the meaning of cooperation. It is important to note that this contribution is subject to two scope conditions. First, joint coordination frameworks are, at best, a second-best solution to the coordination problem. When the first best solution of mutual knowledge is available – i.e., when exchange partners know of all available actions and payoffs and can agree in advance on which actions constitute cooperation – coordination frameworks may play little role. Other solutions to the coordination problem may very well exist. Second and relatedly, since joint frameworks are a solution to the coordination problem, their value is increasing in the complexity of the coordination problem underlying the exchange. When the coordination problem is simple – as is the case when exchange partners can rely on established frameworks (e.g., Uber rides, supermarket shopping), joint frameworks do not contribute much to successful cooperation.

Another significant contribution of this paper is its novel identification strategy – and the experimental platform I developed to test the causal effect of joint frameworks. This identification strategy not only allows us to manipulate participants’ frameworks but also disentangles the coordination problem from exchange partners’ payoffs and the associated risk of defection. Finally, while the experimental results rule in joint frameworks as a solution to the coordination problem, this is but the first study that tests this proposition. Additional experimental and empirical evidence is needed to substantiate the role of joint frameworks and the generalizability of this process working in contexts beyond the shared language modality used here.

Additional Contributions. Beyond elucidating the conditions under which exchange partners can successfully cooperate, my findings contribute to several additional literatures. First, recognizing the role of joint frameworks in successful exchange offers a novel and alternative explanation for the ubiquity of long-term exchange relationships. Scholars have attributed the prevalence of long-term exchange relationships to the fact that it lowers the risk of defection by providing exchange partners with information about each other’s trustworthiness, creates expectations of future cooperative exchange, and because long-term exchanges are often overlaid with interpersonal relationships (Granovetter, 1985; Bull, 1987; Kollock, 1994). Building on this logic, scholars posited that the risk of defection drives actors’ decisions to enter long-term exchange relationships (e.g., Kollock, 1994). My findings suggest that holding the risk of defection constant, long-term relationships are more valuable, and actors are more likely to engage in them when the coordination problem is more complex. In such exchanges, exchange partners have a strong incentive to increase the odds of successful coordination (and thus cooperation) by investing in a joint coordination framework. And, since joint frameworks are

both costly and relationship-specific, exchange partners who have built a joint framework have a strong incentive to continue exchanging with each other (vs. new partners), thus creating a long-term exchange relationship. One implication is that long-term exchange relationships can be valuable even when the risk of defection is negligible. A diner will be much more likely to return to a restaurant that cooked food to her liking than a restaurant that served her an unsatisfying dish, even if both chefs are fully cooperative and the risk of defection is low.

Moreover, Uzzi's (1997: 46) account of a clothing designer's preference to work with long-term manufacturing partners because they better understand his idea of "style" suggests that actors are aware, perhaps intuitively, of the importance of joint frameworks for successful exchange. When that is the case, joint frameworks augment the value of long-term relationships all the more by raising exchange partners' expectations that cooperation will succeed and the relationship will survive into the future. This further increases exchange partners' expectations of future gains from the relationship and thus reduces the risk they will endanger it by defecting. In doing so, this paper complements Axelrod's (1984) analysis of the shadow of the future. Axelrod showed that expectations of future gains from trade deter defection in the present. This paper shows that the light of the past can be just as crucial as the shadow of the future: the past can support cooperation by illuminating how exchange partners should coordinate in both present and future transactions. It is thus not surprising that we observe long-term exchange relationships primarily in exchanges that cannot rely on established frameworks (e.g., restaurants but not cab or Uber drivers; suppliers of customizable parts but not commodities).

Second, my findings also have implications for our understanding of in-group bias. Starting with Tajfel et al. (1971), a long line of experimental work has documented preferential treatment

towards in-group members over out-group members even when group membership was random (known as Minimal Group Experiment or MGE). The presence of in-group bias in MGEs served as a basis for the development of social identity theory (Billig and Tajfel, 1973), theory building on intergroup relations (Messick and Mackie, 1989; Abrams and Hogg, 2006), and other areas of social psychology (Turner et al., 1987). Yamagishi and Kiyonari (2000) found that in-group bias in MGEs is driven by an expectation of generalized reciprocity from in-group members. In-group bias, therefore, exists in simultaneous social dilemma games in which participants must guess whether their partner will cooperate or defect, but not in sequential games in which the partner's choice is known to them.

My findings suggest an additional source for in-group bias – the ability of groups to foster the development of joint coordination frameworks. The traditional MGE and the sequential version of Yamagishi and Kiyonari (2000) have focused on in-group bias in cooperation games that rely on established coordination frameworks. It is thus not surprising that the coordination problem played no role in participants' decisions. Study 2, on the other hand, examines cooperation decisions in sequential cooperation games across two types of games: one with a complex coordination problem and one with a simple coordination problem. Consistent with Yamagishi and Kiyonari (2000), the findings indicate that when the coordination problem is simple, participants do not display a discernable in-group bias in a sequential game. Conversely, when the coordination problem is complex, participants are far more likely to both engage in exchange and successfully cooperate with partners with whom they developed joint frameworks in Study 1. This finding suggests that joint coordination frameworks shared by group members can generate in-group bias even in sequential games when the coordination problem is complex.

Third, recognizing the role of joint coordination frameworks in successful exchange also expands our understanding of the value of closed groups. Social network scholars have long noted that closed groups – that is, groups with dense ties within the group and sparse ties beyond the group – facilitate intra-group cooperation, trust, and norm enforcement (Coleman, 1988; Portes and Sensenbrenner, 1993; for a review see Burt, 2005). My findings offer an additional explanation of the value of closed groups. Closed groups are characterized by repeated interactions between group members that facilitate the creation of shared ways of interpreting the world (Erickson, 1988; Burt, 1992). To the extent these shared interpretations extend to shared understandings of the meaning of cooperation – that is, joint frameworks – they will facilitate in-group cooperation when established frameworks are unavailable, reduce the risk of defection, and generate trust between group members. Consistent with this logic, many canonical examples of closed groups feature groups engaged in exchanges that cannot rely on established frameworks (e.g., Coleman, 1988; Greif, 1989)

In related work, I also explore how joint coordination frameworks facilitate norm enforcement in closed groups by making reputational signals more legible (self-identifying cite). To summarize the argument, third-party norm enforcement necessitates an ability to inform third parties about norm violations through reputational signals. Since these signals are composed of information about who violated a norm, their informativeness relies on a shared understanding of what the norm is and what actions constitute a violation thereof. Study 1 and the long tradition of social-convention-building experiments it builds on suggest that when established frameworks are unavailable, different groups are likely to develop different norms and coordination frameworks. As a result, different groups will have a different understanding of which actions are acceptable. These different understandings will render reputational signals between members

of different groups less informative than in-group reputational signals. Hearing from Anne that Bob has defected on an exchange (violated a norm) is far more informative to Chris if Chris shares Anne's understanding of cooperation and defection (violation). Consequentially, when it comes to exchange and norms that cannot rely on established frameworks, reputational signals will be more informative, and norm enforcement will be stronger within closed groups.

Fourth, the theory and analysis above also contribute to the literature on the embeddedness of economic action. Granovetter (1985) compellingly argues that personal relationships and structures thereof (i.e., networks) generate trust and discourage malfeasance (*ibid*, 490). I augment Granovetter's argument by showing how long-term exchange relationships developed under conditions of complex coordination foster the creation of joint coordination frameworks. These joint frameworks not only increase the odds of successful cooperation by helping exchange partners solve the coordination problem but also increase trust and deter malfeasance by lengthening the shadow of the future.

Fifth, my findings contribute to the game-theoretic literature on cooperation in repeated games. Axelrod (1984) has famously shown that in a repeated prisoner's dilemma game in which coordination is trivial and players can flawlessly execute their strategies, Tit-for-Tat is the dominant strategy. Subsequent analysis has demonstrated that when actors' ability to execute their strategies is prone to error (or "trembling hands"), Tit-for-Tat strategies become inefficient. If actors can erroneously defect when they mean to cooperate, an inevitable accidental defection will induce an endless cycle of retaliation. Molander (1985) found that when actors err, more forgiving strategies that revert to cooperation after several rounds of defection outperform Tit-for-Tat.

Both Axelrod, Molander and the subsequent literature assume the error rate is exogenous. Recognizing the role of coordination frameworks in successful exchange endogenizes the error rate. It suggests that the error rate is a function of the complexity of the coordination problem underlying the exchange. When coordination is easy because established coordination frameworks are available – as is the case in the prisoner’s dilemma game – the error rate is likely to be very low. It is hard to imagine a professional criminal accidentally incriminating a colleague in exchange for leniency. Conversely, when the coordination problem is more complex because exchange partners cannot rely on established frameworks, the error rate will likely be much higher. This is what hindered cooperation in new designer-manufacturer relationships in Uzzi (1997). Finally, exchange partners can endogenize the error rate by building joint coordination frameworks through (successful) joint problem solving. Joint frameworks reduce the complexity of the coordination problem, thus reducing the risk that a well-intentioned action will be perceived as an act of defection and lead to retaliation. Indeed, the failed new designer-manufacturer relationships in Uzzi (1997) would have been more successful had the partners invested in building a joint framework by engaging in low-stakes joint problem solving.

Future Research. The theory and analysis presented above suggest several fertile directions for future research. First, while the experiments presented in this paper constitute important progress on the thorny identification problem, they are best seen as a first cut at a full solution. These experiments employ a common language developed to interpret experimental stimuli as a proxy for a joint coordination framework. While this is a productive approximation commonly used in experimental work on social conventions, it is not a one-to-one mapping to the construct of interest. Our understanding of joint coordination frameworks and their role in cooperation will benefit from more precise measures and experimental manipulations. Second, the findings of this

paper also suggest future directions of inquiry for the study of group identification. Using the contrasting stories of the assimilation (or lack thereof) of 19th-century Chinese immigrants in Jamaica and Guyana, Patterson (1975) argues that individuals choose ethnic and group allegiances based on their interests. When the group does not serve members' interests, group identification declines to the point of assimilation in the public. The finding that in-group bias exists only when group membership benefits actors by generating joint coordination frameworks indirectly supports this assertion. Future research can modify Study 2 to test Patterson's proposition more directly by bookending the study with a questionnaire measuring participants' degree of group identification. If Patterson is correct, we should observe higher degrees of group identification when group membership is more valuable – that is when the coordination problem is more complex, and group members are more likely to develop joint coordination frameworks.

References

Abrams, D., and M. A. Hogg

2006 Social Identifications: A Social Psychology of Intergroup Relations and Group Processes. Routledge.

Almaatouq, A., J. Becker, J. P. Houghton, N. Paton, D. J. Watts, and M. E. Whiting

2021 “Empirica: a virtual lab for high-throughput macro-level experiments.” Behavior Research Methods, 53: 2158–2171.

Axelrod, R.

1984 The Evolution of Cooperation. New York: Basic Books.

Axelrod, R., and W. D. Hamilton

1981 “The Evolution of Cooperation.” Science, 211: 1390–1396.

Azoulay, P., N. P. Repenning, and E. W. Zuckerman

2010 “Nasty, Brutish, and Short: Embeddedness Failure in the Pharmaceutical Industry.” Administrative Science Quarterly, 55: 472–507.

Berinsky, A. J., G. A. Huber, and G. S. Lenz

2012 “Evaluating online labor markets for experimental research: Amazon.com’s mechanical turk.” Political Analysis, 20: 351–368.

Bernstein, L.

2001 “Private commercial law in the cotton industry: Creating cooperation through rules, norms, and institutions.” *Michigan Law Review*, 99: 1724–1790.

Billig, M., and H. Tajfel

1973 “Social categorization and similarity in intergroup behaviour.” *European Journal of Social Psychology*, 3: 27–52.

Bull, C.

1987 “The Existence of Self-Enforcing Implicit Contracts.” *The Quarterly Journal of Economics*, 102: 147–160.

Burt, R. S.

1992 *Structural Holes*. Boston: Harvard Business School Press.

Burt, R. S.

2005 *Brokerage & Closure*. New York: Oxford University Press.

Burt, R. S., R. E. Reagans, and H. C. Volvovsky

2021 “Network Brokerage and The Perception of Leadership.” *Social Networks*, 65: 33–50.

Canales, R.

2014 “Weaving Straw into Gold: Managing Organizational Tensions Between Standardization and Flexibility in Microfinance.” *Organization Science*, 25: 1–28.

Canales, R., and J. Greenberg

2016 “A Matter of (Relational) Style: Loan Officer Consistency and Exchange Continuity in Microfinance.” *Management Science*, 62: 1202–1224.

Centola, D., and A. Baronchelli

2015 “The spontaneous emergence of conventions: An experimental study of cultural evolution.” *Proceedings of the National Academy of Sciences*, 112: 1989–1994.

Coleman, J. S.

1988 “Social Capital in the Creation of Human Capital.” *American Journal of Sociology*, 94: S95–S120.

DiBenigno, J.

2020 “Rapid Relationality: How Peripheral Experts Build a Foundation for Influence with Line Managers.” *Administrative Science Quarterly*, 65: 20–60.

Elfenbein, D. W., and T. R. Zenger

2014 “What Is a Relationship Worth? Repeated Exchange and the Development and Deployment of Relational Capital.” *Organization Science*, 25: 222–244.

Erickson, B. H.

1988 “The Relational Basis of Attitudes.” In B. Wellman, S. D. Berkowitz, and M. S. Granovetter (eds.), *Social Structures: A Network Approach*: 99–121. Cambridge, UK: Cambridge University Press.

Gibbons, R. S., and R. M. Henderson

2012 “Relational Contracts and Organizational Capabilities.” *Organization Science*, 23: 1350–1364.

Granovetter, M. S.

1985 “Economic Action and Social Structure: The Problem of Embeddedness.” *American Journal of Sociology*, 91: 481–511.

Greif, A.

1989 “Reputation and Coalitions in Medieval Trade: Evidence on the Maghribi Traders.” *Journal of Economic History*.

Guilbeault, D., A. Baronchelli, and D. Centola

2021 “Experimental evidence for scale-induced category convergence across populations.” *Nature Communications*, 12: 1–7.

Gulati, R.

1995 “Does Familiarity Breed Trust? The Implications of Repeated Ties for Contractual Choice in Alliances.” *Academy of Management Journal*, 38: 85–112.

Gulati, R., F. Wohlgezogen, and P. I. Zhelyazkov

2012 “The Two Facets of Collaboration: Cooperation and Coordination in Strategic Alliances.” *Academy of Management Annals*, 6: 531–583.

Hansen, M. T.

1999 “The Search-Transfer Problem: The Role of Weak Ties in Sharing Knowledge across Organization Subunits.” *Administrative Science Quarterly*, 44: 82–111.

Hart, O. D., and J. Moore

1999 “Foundations of incomplete contracts.” *Review of Economic Studies*, 66: 115–138.

Kollock, P.

1994 “The emergence of exchange structures: An experimental study of uncertainty, commitment, and trust.” *American Journal of Sociology*, 100: 313–345.

Kreps, D. M.

1990 “Corporate Culture and Economic Theory.” In Alt J. and Shepsle K. (ed.), *Perspectives on Positive Political Economy*: 90–143. Cambridge, MA: Cambridge University Press.

Lawler, E. J., S. R. Thye, and J. Yoon

2009 *Social Commitments in a Depersonalized World*. Russel Sage Foundation.

Messick, D. M., and D. M. Mackie

1989 "Intergroup Relations." *Annual Review of Psychology*, 40: 45–81.

Molander, P.

1985 "The Optimal Level of Generosity in a Selfish, Uncertain Environment." *Journal of Conflict Resolution*, 29: 611–618.

Molm, L. D., D. R. Schaefer, and J. L. Collett

2009 "Fragile and resilient trust: Risk and uncertainty in negotiated and reciprocal exchange." *Sociological Theory*, 27: 1–32.

Paolacci, G., and J. Chandler

2014 "Inside the Turk: Understanding Mechanical Turk as a Participant Pool." *Current Directions in Psychological Science*, 23: 184–188.

Patterson, O.

1975 "Context and Choice in Ethnic Allegiance: A Theoretical Framework and Caribbean Case Study." In N. Glazer and D. P. Moynihan (eds.), *Ethnicity: Theory and Experience*. Harvard University Press.

Portes, A., and J. Sensenbrenner

1993 "Embeddedness and Immigration: Notes on the Social Determinants of Economic Action." *American Journal of Sociology*, 98: 1320–1350.

Powell, W. W.

1990 "Neither market nor hierarchy." *Research on Organizational Behavior*, 12: 295–336.

Reagans, R. E., and B. McEvily

2003 "Network Structure and Knowledge Transfer: The Effects of Cohesion and Range."

Administrative Science Quarterly, 48: 240–267.

Schelling, T. C.

1960 *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.

Simpson, B., and R. Willer

2015 "Beyond Altruism: Sociological Foundations of Cooperation and Prosocial Behavior."

Annual Review of Sociology, 41: 43–63.

Tajfel, H., M. Billig, R. Bundy, and C. Flament

1971 "Social Categorization in Intergroup Behaviour." *European Journal of Social*

Psychology, 1: 149–178.

Turner, J. C., M. A. Hogg, P. J. Oakes, S. D. Reicher, and M. S. Wetherell

1987 *Rediscovering the Social Group: A Self-Categorization Theory*. Oxford: Blackwell.

Uzzi, B.

1997 "Social Structure and Competition in Interfirm Networks: The Paradox of

Embeddedness." *Administrative Science Quarterly*, 42: 35.

Weber, R. A., and C. F. Camerer

2003 “Cultural Conflict and Merger Failure: An Experimental Approach.” *Management Science*, 49: 400–415.

Wilkes-Gibbs, D., and H. H. Clark

1992 “Coordinating beliefs in conversation.” *Journal of Memory and Language*, 31: 183–194.

Yamagishi, T., and T. Kiyonari

2000 “The Group as the Container of Generalized Reciprocity.” *Social Psychology Quarterly*, 63: 116.

Zuckerman Sivan, E. W.

2017 *Why Identity? A Prolegomenon to any Account of Social Organization or Human Action.*

Figure 1: Coordination Game




Figure 2: Images in the Established Frameworks: Available Condition




Figure 3: Images in the Established Frameworks: Unavailable Condition




Figure 4: Principal's Screen when Deciding to Hire Agent




Figure 5: Principal’s Screen when Describing Image to Agent




Figure 6: Agent’s Screen when Deciding Whether to Pay to Read Principal’s Message




Figure 7: Agent's Screen when Selecting an Image




Figure 8: Principal's Screen when Observing the Outcome




Figure 9: Cooperation Game Chart




Figure 10: Distribution of Successful Rounds




Figure 11: Image from the Established Frameworks: Unavailable Condition




Figure 12: Image from the Established Frameworks: Unavailable Condition




Figure 13: Image from the Established Frameworks: Unavailable Condition




Figure 14: Successful Rounds by Condition




Table 1: Severity of the Coordination Problem

Severity of the Coordination Problem		Established Coordination Frameworks	
		Available	Unavailable
Joint frameworks	Yes	Low	Medium
	No	Low	High

Table 2: Experimental Conditions

Experimental Conditions		Established Coordination Frameworks	
		Available	Unavailable
Engaged in Exchange	Yes	I	II
Involving Low-Stakes Joint Problem Solving	No	III	IV

Table 3: The Meaning of Cooperation and Defection

Principal's Point of View		Agent's Point of View		
	Cooperation	Defection	Cooperation	Defection
Principal	Hire Agent (-50)	Terminate (0)	Hire Agent (+50)	Terminate (0)
Agent	Pick Correct Image (+80)	Pick Wrong Image (0)	Pay to Read Principal's Message (-30)	Don't Pay to Read Principal's Message (0)

Table 4: Participant Demographics

Experimental Condition	Count	Age		Gender			Education			
		Avg.	Std.	Male	Female	Other/NA	HS	BA	MS or Higher	Other/NA
Total	1,070	38.9	11.7	45.1%	46.5%	8.4%	31.1%	41.6%	16.2%	11.1%
I	224	39.4	11.6	42.9%	50.1%	7.0%	32.6%	43.3%	18.3%	5.8%
II	224	39.6	12.0	47.3%	44.2%	8.5%	33.5%	38.9%	14.7%	12.9%
III	306	38.4	12.1	48.0%	44.8%	7.2%	30.1%	44.1%	15.4%	10.4%
IV	316	38.6	11.2	42.4%	46.8%	10.8%	29.4%	39.9%	16.5%	14.2%
Ethnicity										
	Caucasian		African American		Asian	Hispanic	Other/NA			
Total	74.6%		6.0%		5.2%	4.6%	9.6%			
I	77.2%		4.9%		8%	4.5%	5.4%			
II	73.7%		7.6%		2.7%	4%	12%			
III	72.5%		9.2%		3.6%	4.9%	9.8%			
IV	75.3%		2.5%		6.6%	4.7%	11.9%			

Table 5: OLS Regression Results

	(1)	(2)	(3)	(4)
VARIABLES	Successful Cooperation	Rounds Played	Effort Rate	Success Rate
Joint Problem Solving	0.250 (0.249)	0.232 (0.250)	0.0526 (0.0368)	0.0500 (0.0497)
Established Frameworks: Unavailable	-1.982*** (0.249)	-1.179*** (0.250)	-0.0682* (0.0375)	-0.396*** (0.0497)
Joint Problem Solving x Established Frameworks: Unavailable	1.250*** (0.352)	0.893** (0.354)	0.0120 (0.0524)	0.250*** (0.0703)
Constant	3.741*** (0.176)	3.955*** (0.177)	0.899*** (0.0260)	0.748*** (0.0352)
Observations	448	448	382	448
R-squared	0.179	0.079	0.027	0.179

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1