The Dynamics of Strategic Alliances: Theory and Experimental Evidence

Albert Banal-Estañol* Tobias Kretschmer†

Debrah Meloso ‡ Jo Seldeslachts§

November 2015

Abstract

Cooperation and coordination problems during a strategic alliance's implementation are considered two key reasons why alliances fail. We propose a framework of post-formation alliance dynamics to identify the causes for cooperation and coordination problems and their subsequent impact on the dynamic performance of inter-firm alliances. We show that the degree of partners' complementarities and the level of product market competition between partners interact in determining if an alliance faces cooperation or coordination issues. We then experimentally study the performance of cooperation and coordination alliances and find that coordination alliances outperform cooperation alliances. Interestingly, their better performance is predominantly due to partners collaborating from the start more frequently. Differences in the maintenance of successful and recovery of unsuccessful alliances are much less pronounced.

*Department of Economics and Business, Universitat Pompeu Fabra and Department of Economics, City University of London. E-mail: albert.banalestanol@upf.edu. Web: http://albertbanalestanol.com

†Institute for Strategy, Technology and Organization and Organizations Research Group, LMU Munich. E-mail: t.kretschmer@lmu.de

‡Dpt. of Accounting and Finance, ESC-Rennes School of Business. E-mail: debrah.meloso@esc-rennes.com

§DIW Berlin, KU Leuven and University of Amsterdam. E-mail: jseldeslachts@diw.de

1 Introduction

Strategic alliances, broadly defined as interfirm cooperative arrangements aimed at pursuing mutual strategic objectives, are increasingly used to pursue strategic aims that would be beyond the capabilities of a single firm (Gulati, Wohlgezogen, and Zhelyazkov 2012; Garrette, Castañer, and Dussauge 2009). Indeed, strategic alliances have become such a frequent strategic choice that firms, rather than managing individual alliances, frequently manage their alliance portfolios (Wassmer and Dussauge 2011). Despite their frequency, success rates remain surprisingly low, even for individual alliances. Failure rates are often well in excess of 50% (Kale, Dyer, and Singh 2002; Lunnan and Haugland 2008). Hence, alliances, despite their potential, often produce disappointing outcomes (Kale and Singh 2009; Kogut 1989; Park and Ungson 1997; Park and Russo 1996; Reuer and Zollo 2005).

The fact that so many alliances fail to meet their intended targets suggests that there are significant impediments to alliance success (Anand and Khanna 2000; Gottschlag and Zollo 2007). Broadly speaking, we distinguish three phases to identify and study factors that cause failure: the pre-formation phase that sets the initial conditions and determines early interactions, the continuation phase that determines the degree to which ongoing benefits are realized, and finally an alliance resilience phase, where alliance partners try to reestablish collaboration following a setback.

In the pre-formation phase, the characteristics of the alliance partners and their relationship matter most as there is less alliance history to consider. Specifically, close rivals may experience tensions when forming an alliance (Kale, Singh, and Perlmutter 2000; Khanna, Gulati, and Nohria 1998; Park and Russo 1996). Conversely, strong complementarities between partnering firms would align incentives and carry more scope for success (Gulati and Sytch 2007; Silverman and Baum 2002). Problems in the alliance's post-formation stage, due to e.g. a lack of trust or wrong actions undertaken in the early stages of the alliance, are often considered as the main factor in explaining failures (Park and Ungson 2001; Uzzi 1997; Zaheer, McEvily, and Perrone 1998). Interestingly, the resilience of alliances, i.e. the ability to bounce back, has not yet been covered in prior research.

Importantly, the initial characteristics of an alliance are likely to have repercussions for the way alliances are played out over time. Put differently, the incentives and mechanisms present in an alliance, either by design or exogenously, will play an important role in the success of an alliance. However, these initial conditions can manifest in several ways: first, they may lead to differences in the initial level of cooperation. Second, different alliances may prove more or less stable in their behavior over time. And third, as alliances often have to deal with setbacks, different alliance types can display different behavior in response to a (temporary) breakdown in cooperation.

We provide a framework capturing all these elements of an alliance to study in detail the drivers of success or failure over time. We capture both the pre-formation phase to characterize which characteristics specific to a set of two partners hwill generate particular types of alliance, and the post-alliance formation phase to characterize post-formation dynamics in line with Doz and Hamel (1998, p. XV), who state "No matter how carefully organizations strategize about partner selection and initial governance setup, many scholars believe that managing the alliance relationship over time is usually more important."

Our core distinction is between *cooperation* and *coordination* (Gulati, Wohlgezogen, and Zhelyazkov 2012; Park and Ungson 2001). *Cooperation* refers to issues where individual incentives may not be aligned with maximizing the value of the alliance, while *coordination* captures alliance partners' need to settle on a mutually beneficial solution that requires collaborative behavior by all alliance partners. We start by identifying the factors contributing to alliances being of one or the other type in the pre-formation stage and the consequences for performance and stability in the alliance implementation stage.

We first propose a formal model in which we identify important parameters that will determine the type of alliance, specifically the degree of rivalry in the product market and the degree of complementarities within an alliance. Second, to assess the dynamic properties and the respective alliance types' resilience to breakdown, we conduct a series of experiments resembling the theoretical setup, but including an option to terminate the alliance. Thus, the studies closest to ours in both research idea and methodology are Arend (2009); Amaldoss and Staelin (2010), and Agarwal, Croson, and Mahoney (2010). Arend (2009) experimentally tests a repeated Prisoner's Dilemma alliance game with the option to exit. Amaldoss and Staelin (2010), among other things, experimentally compares cross-function and same-function alliances, where based on surveys of managers they postulate that contributions into cross-function alliances are multiplicative and contributions into same-function alliances are additive. Finally, Agarwal, Croson, and Mahoney (2010) experimentally compare alliances with high and low common benefits for its members, relative to private benefits.

Our paper contributes to a growing literature in strategy and organizations employing mixed methods to highlight multiple important issues of a question. In our case, the benefit of deriving an alliance typology from its microfoundations and two fundamental characteristics of alliances (i.e. potential complementarities and rivaly outside the alliance) and testing such a model experimentally derives mainly from the ability to observe alliance partners' behavior in the lab while being able to match this to structural parameters derived from a fully rational model. Where static alliance games are unrealistic and dynamic ones can generate almost any behavior in equilibrium, our experimental setup lets us focus on the behavior we can observe.

Our results suggest that coordination alliances perform better than cooperation alliances in all dimensions: better actions are undertaken, first-best outcomes are achieved more often and successful alliances last longer. We find that this performance advantage of coordination alliances originates largely from the initial conditions: participants in cooperation alliances fail to collaborate more often in the first periods or do not start in the first place. Starting from this higher level of collaboration, coordination alliances then also maintain collaboration with a higher likelihood, although the difference is not as stark as in the initial play. Moreover, following a breakdown in collaboration, partners in cooperation alliances will reverse to collaboration less frequently than those in coordination alliances. In sum, our results indicate that the dynamics of well-functioning alliances are not so different across alliance types. However, partners in cooperation alliances fail more often to collaborate from the start and have more trouble to turn bad situations around.

2 Theoretical model

2.1 Cooperation and Coordination as Problems in Strategic Alliances

Many studies assert that problems of cooperation, stemming from divergent incentives between alliance partners, are the root cause of many failed alliances (Kogut 1988; Oxley 1997; Sampson 2004; Arend and Seale 2005; Parkhe 1993; Gulati 1995; Robinson and Stuart 2007). Conflicting interests can cause diminished commitment of partners over time that erodes the relationship (Doz 1996). Even worse, they can lead to opportunistic behavior from the start and to the pursuit of self-interest with no regard for unenforceable commitments (Williamson 1985). However, an alternative view suggests that problems of coordination, i.e. the failure to align partners' actions, can contribute significantly to underperformance (Gulati, Wohlgezogen, and Zhelyazkov 2012; Gulati, Lawrence, and Puranam 2005; Gulati and Singh 1998). Indeed, alliances involve at least some division of labor and there is usually some task interdependence among partners. This creates coordination challenges, since a precise synchronization of activities is difficult and failure to do so is costly (Varshney and Oppenheim 2011). During implementation, then, a lack of alignment between partners' actions may become apparent and coordination failures occur.

Yet, although recent studies highlight the importance of both problems in explaining failures, it is less clear *when* and *how* cooperation concerns dominate coordination concerns (or vice versa), and what this means for the subsequent evolution of a partnership. Accordingly, we study in detail both the causes and consequences of problems of cooperation and coordination during an alliance's post-formation phase.

2.2 A Model of Alliance Types

To this end, we propose a formal (game-theoretical) model of the various types of strategic interaction within alliances. Our model describes which partner characteristics lead to the pervasive problems of cooperation and coordination. Specifically, after describing the general setup, we specify the partners' profits and how they are affected by the decisions made within the alliance. Based on the resulting (static, i.e. one-shot) profits, we characterize the different types of strategic interaction and equilibrium outcomes.

General setup We focus on alliances where their potential synergies are only realized when both partners contribute resources (e.g., Hagedoorn 1993; Agarwal, Croson, and Mahoney 2010). Thus, alliance actions are interdependent in the sense that a mutual engagement of members is required to reach a partnership's full value (Dyer 1996; Dyer and Singh 1998).[1] Moreover, while each partner decides how much to contribute towards the alliance activity (e.g. how much time and effort of key personnel to dedicate to the partnership), the rewards of their individual contributions are common to all alliance partners (Agarwal et al. 2012). In our model, we focus on alliances that aim to reduce joint production costs, as in information sharing or R&D joint ventures (JVs). However, our setup can

[1] Alliances that are not set up for synergies do not require interdependent actions to prosper. Therefore, these partnerships are by definition less likely to encounter problems of coordination or cooperation, and hence should not lead to failures due to wrong actions undertaken by its members. See e.g. Oxley (1997) for a detailed typology of alliances.

also be interpreted as an alliance where partners collaborate to increase aggregate market demand, as in marketing JVs. Alliance members, though, remain independent firms, and might even be close competitors in the product market.

Profit specification Two symmetric firms, distinguished by the index $i = 1, 2$ when necessary, have formed a strategic alliance. The profits of each partner, $\pi(\tilde{c}, d)$, depend on a parameter $\tilde{c}$ affecting the costs of production, as well as on a parameter d affecting the intensity of competition between the companies in the product market. We use a fairly general specification of profits but we impose certain restrictions. As it is natural to assume, profits decrease in the level of costs as well as in the intensity of competition:

$$\frac{\partial \pi(\tilde{c}, d)}{\partial \tilde{c}} < 0 \text{ and } \frac{\partial \pi(\tilde{c}, d)}{\partial d} < 0. \tag{1}$$

In addition, we assume that the reduction in profits due to higher costs is less important as the level of competition in the product market increases, i.e.

$$\frac{\partial^2 \pi(\tilde{c}, d)}{\partial d \partial \tilde{c}} > 0. \tag{2}$$

This relationship between costs and competition can be traced back to the *Schumpeterian view* that competition lowers the benefits of investment in cost-reducing innovations. In terms of standard economic models, this condition holds, for example, in the classic linear oligopoly models of Bertrand-price and Cournot-quantity competition. Indeed, as shown in Appendix A, if $\tilde{c}$ represents the constant marginal costs of production and d the degree of substitutability between the products of the two firms, then the cross-partial derivative of the equilibrium profits is positive, as our condition 2 stipulates.[2]

Resource contribution When partners decide to contribute resources to the alliance, production costs decrease as a result. Modeling decisions as a binary choice, each firm's cost of production, $\tilde{c}$, is equal to c if no partner contributes resources; reduced to $c - r$, with $r \in (0, c)$, if only one firm contributes resources; and to $c - kr$, $k \in (1, c/r)$, if they both contribute resources, where the parameter k represents the degree of complementarity

[2]More generally, this condition holds when a more competitive product market leads to each firm facing a lower individual demand (Vives 2008) or when intellectual property rights do not perfectly protect firms' investments/innovations (Gilbert 2006). See also Rafique (2013) for empirical evidence of a negative relationship between competition and innovation.

between partners' resources. The more partners are specialized in different technologies or the more dissimilar resources they bring into the alliance, the more complementarities k are present.[3] Through this, partners' actions become more interdependent: a commitment of both leads to better outcomes in the alliance.

Note that the cost of production $\tilde{c}$ of *each* partner is lowered when resources are contributed. Thus, partners benefit from each other's contribution. Indeed, the economic value created in the alliance is mostly commonly shared, as it is difficult to exclude alliance partners from enjoying these gains, regardless of whether or not they contributed toward the economic value creation (Gulati and Sytch 2007; Agarwal, Croson, and Mahoney 2010).

Moreover, we take that the extra gains from contributing resources are greater if the other party also contributes resources. We follow Dyer and Singh's (1998, p. 666) concept of *strategic complementarities*, which are "complementary resource endowments from distinctive resources of alliance partners that collectively generate greater rents than the sum of those obtained from the individual endowments of each partner." Complementary resources have been discussed widely as a key factor in driving synergistic returns from alliances (see e.g., Hamel 1991; Harrigan 1985; Hill and Hellriegel 1994; Teece 1988). Although the complementary resources create the potential for rents, these gains can only be realized if both partners contribute resources. Indeed, a primary reason for failure in alliances is not that the two partners do not possess the right resources, but rather because they do not choose the right actions with these resources inside the alliance. In formal terms,

$$\pi\left(c-kr,d\right)-\pi\left(c-r,d\right)>\pi\left(c-r,d\right)-\pi\left(c,d\right). \tag{3}$$

As shown by this inequality, alliance benefits can be divided in two parts: gains through individual contributions, reducing costs from c to $c-r$ (given by the right-hand side), and synergistic gains through joint efforts, reducing costs from $c-r$ to $c-kr$ as opposed to from c to $c-r$ (given by the difference between the left-hand side and the right-hand side).

[3] As an example of a partnership with high complementarities, labeled as a *cross-function* alliance in their study, Amaldoss and Staelin (2010) mention the collaboration set up in 2001 by IBM, Sony, and Toshiba to develop high-performance, low-power semiconductors for consumer electronics. IBM brought its knowledge of chip technology, Sony its expertise in consumer electronics, and Toshiba its manufacturing knowledge. In contrast, Advanced Micro Devices (AMD) and Fujitsu pooled their expertise in flash memory chips to jointly produce such chips. In this alliance, partners pool similar resources and complementarities are thus low (*same-function* alliances in Amaldoss and Staelin's terminology). Our typology of alliances in terms of the degree of complementarities is further consistent with Park and Russo's (1996) distinction between *integrative* and *sequential* alliances, and Porter and Fuller's (1986) *X-form* and *Y-form* coalitions.

Clearly, the larger the degree of complementarity k between partners, the higher the extra rents that can be generated through joint efforts.

While joint efforts lead to higher benefits, partners' interests are not completely aligned. Alliance partners remain independent actors and retain control over their own resource contributions (Deeds and Hill 1996; Hamel 1991; Park and Ungson 2001). We include this feature in our model by assuming that each partner's resource contribution requires a private cost, equivalent to e monetary units. Some of these resource contributions can be contracted upon, but enforcing (near-)complete contracts would be prohibitively expensive (Crocker and Reynolds 1993). For simplicity, we assume resource contributions to be fully non-contractible.[4]

Table 1 summarizes the profits of each firm as a function of both firms' choices within the alliance. We denote by C the decision to contribute resources and, similarly, by nC the decision not to do so. We further assume that contribution decisions are taken simultaneously, or in other words, that alliance members cannot condition behavior on the decisions of their partner. Our model reflects thus the uncertainties resulting from difficulties in monitoring the behavior of alliance partners (e.g., Agarwal, Croson, and Mahoney 2010). Indeed, many of the required actions are simply too difficult to observe and describe in sufficient detail and are often plagued by ambiguity. Action interdependencies make it even more difficult to measure separate contributions immediately (Gulati and Singh 1998; Mesquita, Anand, and Brush 2008).[5]

[Table 1 about here.]

Strategic behavior within alliances The following two propositions (proofs in Appendix B) state the essential features of the strategic interaction between alliance partners, and the resulting Nash equilibria, as a function of the *degree of competition between alliance members* outside the alliance, d, and the *degree of complementarity of partners' resources* inside the alliance, k.

Proposition 1. Unambiguous Alliances

There exist critical degrees of competition, d^*, and complementarities between partners,

[4] If all contributions were contractible, efficient decisions would always be made in an alliance, which is clearly inconsistent with what we observe in reality.

[5] This is especially true for technology alliances, where tacit – and hence unobservable and uncontractible – knowledge typically plays a big role (Oxley 1997).

$k^*(d)$ and $k^{**}(d)$, with $\partial k^*/\partial d > 0$ and $\partial k^{**}/\partial d > 0$, such that alliances have:
(i) *Clear commitment*: for $d \leq d^*$, both partners contributing resources is an equilibrium in dominant strategies. This equilibrium Pareto-dominates no one contributing.
(ii) *Clear non-commitment*: for $d > d^*$ and $k \leq k^{**}(d)$, no one contributing is an equilibrium in dominant strategies. This equilibrium Pareto-dominates both contributing.

If the degree of market competition between alliance partners is sufficiently low (case (i)), the resulting strategic interaction in the alliance poses no problems: both partners contribute resources since it is, independently of what the partner does, the best action to undertake. Similarly, if competition is strong and the degree of complementarities is sufficiently low (case (ii)), no partner contributes resources and this is the efficient outcome.

These are what we call clear commitment and clear non-commitment alliances, since the optimal actions for each partner are independent of what the other does and the results always generate the best possible, i.e. they are Pareto dominant.[6] While both types of situations surely occur often, they do not pose any strategic issue, i.e. they have unique and efficient equilibria. Therefore, clear commitment and clear non-commitment alliances will not be discussed further. The following proposition then represents types of alliances that represent potential problems.

Proposition 2. Alliances with Coordination or Cooperation Problems
There exist critical degrees of competition, d^*, and complementarities between partners, $k^*(d)$ and $k^{**}(d)$, with $\partial k^*/\partial d > 0$ and $\partial k^{**}/\partial d > 0$, such that alliances have:
(i) *Coordination problems*: for $d > d^*$ and $k > k^*(d)$, there are two equilibria; both or no one contributing resources, whereby the former Pareto-dominates the latter.
(ii) *Cooperation problems*: for $d > d^*$ and $k^{**}(d) < k \leq k^*(d)$, no one contributing is an equilibrium in dominant strategies, but it is Pareto-dominated by both contributing.

If the level of competition is high and the degree of complementarities relatively high, inefficient results become possible. If complementarities are high relative to the degree of product market competition (case (i)), partners interact as in the standard "Coordination Game" (see e.g., Gibbons 1992). Although there is no real conflict of interest, partners

[6] We use thus the concept of *Pareto dominance*: a situation is Pareto dominant or Pareto efficient if it is impossible to make a partner better off without making the other partner worse off (see e.g., Fudenberg and Tirole 1984).

do not automatically coordinate on the good result where both contribute. This is due to the fact that a partner only wants to contribute if it believes the other does so as well. Intuitively, high complementarities lead to relatively high payoffs when contributing resources, but *only* if the other contributes as well.[7]

Instead, if complementarities are low for the degree of product market competition (case (ii)), partners will not contribute. Indeed, not contributing becomes a dominant strategy. This is because as competition becomes stronger, the gains from not contributing resources when the other has done so become more important, as this gives an important extra competitive edge in the market where both firms are competing. Thus, partners end up not contributing but it would have been efficient to do so. The alliance resembles a "Prisoners Dilemma". In other words, while both choose not to contribute resources, it would have been beneficial for each partner to contribute.

As an example, Figure 1 illustrates all alliance types for the classic linear oligopoly model of Cournot-quantity competition. Clear commitment alliances correspond to the dark area below the horizontal line ($d < d^*$). Alliances with coordination problems correspond to the dark area above the horizontal line ($d > d^*$ and $k > k^*(d)$). Alliances with cooperation problems correspond to the intermediate area above the horizontal line ($d > d^*$ and $k^{**}(d) > k > k^*(d)$). Finally, alliances with clear non-commitment correspond to the lighter area above the horizontal line ($d > d^*$ and $k < k^{**}(d)$).[8]

[Figure 1 about here.]

Discussion We now embed our resulting alliance types in the literature by comparing our non-obvious alliances outlined in Proposition 2 with related conceptual papers.

First, our framework identifies alliances with problems of coordination for relatively high complementarities and low competition (case (i) of Proposition 2). There are two recent studies that identify alliances with similar characteristics. As mentioned before, Amaldoss and Staelin (2010) relate high complementarities with cross-function alliances,

[7] Note that we use the term coordination failure to denote situations in which partners could have achieved a better equilibrium had they better coordinated their provision of resources. Conversely, the organizational literature on coordination failure often focuses on problems that affect the combination or integration of resources in a joint effort (Gulati, Wohlgezogen, and Zhelyazkov 2012). This conception of coordination failure, thus, concentrates on these combination and integration problems. We hope to provide a complementary view with our study.

[8] In the light area below the horizontal line, resource contributions are not strategic complements, and hence not part of our framework.

where partners bring specialized resources into the collaboration. Based on surveys of managers, these alliances are then postulated to be characterized by a multiplicative function, i.e. individual efforts into the alliance only yield benefits if all partners contribute. This characterization naturally leads to several equilibria, where all partners contributing Pareto dominates no one contributing. Based on the work by Khanna, Gulati, and Nohria (1998), Agarwal, Croson, and Mahoney (2010) define alliances in terms of generating private benefits – accruing only to individual firms – versus generating common benefits – accruing to all partners. Having relatively higher private benefits puts partners in a more "competitive" interaction akin to a higher degree of product market competition and lower complementarities in our setup. Indeed, in a numerical example, they show that when common benefits are relatively high with respect to private benefits, then all partners contributing and no one contributing are equilibria, where the first is Pareto dominant.

Second, our framework identifies alliances with problems of cooperation for relatively low complementarities with respect to competition (case (iii)). Similarly, Agarwal, Croson, and Mahoney (2010) show in a numerical example that when common benefits are relatively low with respect to private benefits, then no partner contributing is the only equilibrium.

While the abovementioned papers clearly have elements in common with our setup, i.e. they make a link between partner characteristics and alliance interactions, we provide a formal reasoning of why and how these connections work.[9] Indeed, our game-theoretical model has linked two key partner characteristics to their choices when deciding whether to contribute resources into the alliance. We have shown that the degree of complementarity and the level of product market competition interact in determining if and when an alliance shows problems of coordination or cooperation.

3 Experimental Design

We now describe the experimental design we use to study the consequences of the problems of coordination and cooperation within alliances. This lets us extend our setup to allow for dynamic interactions and the option to exit the alliance. Indeed, alliance partners interact

[9] There is a series of studies that frame alliances as a game with Coordination Games or Prisoner 's Dilemma, but without making an explicit link to partner characteristics. Notably, Gulati (1995) was the first to point out that while alliances are often viewed as Prisoner's Dilemma games, they should often be more appropriately characterized by a coordination game. Most often, however, conceptual papers on alliances that apply game theory just assume a Prisoner's Dilemma to be the best characterization (e.g., Arend 2009; Parkhe 1993; Parkhe 1993; Zeng and Chen 2003).

repeatedly, can change their decisions over time and react to partner's past choices. For example, a partner might start contributing resources but decide to stop doing so if her partner does not collaborate, or even decide to terminate the alliance. Thus, apart from making our setup dynamic, we also think it is important to allow for the decision to terminate the alliance, as in reality firms have this choice.[10] We first introduce the exit option and describe the matrix of payoffs that alliance partners in our experimental game can get in each period of time. Subsequently, we introduce dynamics and define alliance success and failure, based on performance relative to alternative alliance outcomes.

Exit option In each period, we assume that the alliance is created or maintained (if already started in a previous period) if and only if both firms agree to do so. If any of the two firms refuses to participate, both of them obtain the profits of the outside exit option ("E"), which are equal to the profits of no resource contribution, $\pi(c, d)$. When both partners agree to be in the alliance, they then have to decide whether to contribute non-verifiable resources at a cost e ("C" or "nC"), as described in the previous section. In addition, keeping an alliance going involves a fixed and observable cost F to each partner, attributable to, e.g. buildings, machines, administration and other overhead that is needed to keep the alliance in place. Table 2 describes the profits from each combination of actions. We take F to be positive but relatively small, so that both participating and not contributing resources is worse than not participating, but both participating is better if partners subsequently contribute resources.[11]

[Table 2 about here.]

Per-period payoff matrix We ask experimental subjects to play an "alliance" game with problems of coordination or cooperation. As a representative case for the alliance with coordination problems, which we will call "CD-alliance" from now on, we use the per-period payoff matrix described in panel (a) of Table 3. For the alliance with cooperation

[10]Moreover, allowing for this option makes it easier to compare our results with empirical papers; indeed, most of the empirical literature identifies an alliance's performance through its duration, as other measures are often unavailable. Perhaps surprisingly, we are aware of only one other theoretical/experimental paper on alliances that allows for firms to decide to be in or out of the alliance. Based on Arend and Seale (2005), Arend (2009) models and experimentally tests an alliance as an iterated Prisoner's Dilemma game with the option to exit.

[11]In technical terms, if $\pi(c - kr, d) - e > \pi(c - r, d)$ then $\pi(c - kr, d) - e - F > \pi(c, d)$. For a large F, participation in the alliance is never beneficial, and thus the alliance is never started in the first place.

problems, "CP-alliance", we use the payoff matrix of panel (b). These matrices have the same payoff orderings and strategic properties as the alliances with coordination and cooperation problems from the theoretical framework section and include a small fixed cost of participating in the alliance as described in the previous subsection (Table 2). But we do not use the payoffs that would arise from a particular parameter specification of a particular model of competition for two reasons. First, we prefer to use single-digit payoffs and avoid unnecessarily complex ciphers and decimals. Thus, we make the setup easy to implement in the laboratory. Second, we choose two matrices with minimal payoff differences – as if we were taking a "regression discontinuity" approach – to isolate the effects of the change in strategic interaction.[12] Indeed, the payoff matrices are exactly identical, except for the payoff of outcome $[nC, C]$, which is one unit larger than the payoff of outcome $[C, C]$ in the CP-alliance and one unit smaller in the CD-alliance.

[Table 3 about here.]

Note that if there was only one period, the (non-weakly dominated) Nash equilibria would be as follows. In the CP-alliance, both firms would decide not to participate because they anticipate that they will not contribute within the alliance. Conversely, in the CD-alliance, they might either participate and contribute resources therein or not participate because they anticipate that they would not contribute resources therein.[13]

Dynamic modeling We model dynamics by considering an infinite repetition of the decisions and payoffs described in the previous subsection, while including a per-period common discount factor. In each period, alliance partners decide whether to set up or maintain the alliance, and if both choose to do so, they each decide whether to contribute resources. Participants maximize the discounted sum of stream of profits, which is equivalent to firms maximizing their expected utility and facing an exogenous probability of

[12]Although in an experiment a regression discontinuity design is not used to establish causality, its intuition is similar to the standard econometric use of the term (see e.g., Hahn, Todd, and Klaauw 2001). One can minimize other factors – different from the strategic awareness in which we are interested – by considering circumstances that are as similar as possible, but lying on opposite sides of the "treatment divide"(i.e., problems of cooperation versus problems of coordination).

[13]Both firms taking the outside option, while they anticipate that participating gives them larger profits would be an equilibrium in weakly dominated strategies. This is an equilibrium because if one of them stays out, the decision of the other does not affect the profits. Forward induction arguments could help in selecting the best equilibrium in this scenario. Indeed, the participation decision of one firm could signal that it will contribute because, if it will not, it would earn less than with its outside option. The (costly) participation decision might help in coordinating on the best outcome; see e.g. Cachon and Camerer (1996) for a detailed exposition of this situation.

termination in each period (see e.g., Arend 2009; Agarwal, Croson, and Mahoney 2010). In our setup, we induce a discount factor of $\delta = 0.9$ by including an exogenous probability of termination of 0.1 at the end of each period.

We modify this infinitely repeated setup slightly. First, we assume that if at least one partner refuses to maintain the alliance in a given period, then the alliance ends and cannot be restarted in a subsequent period. In this case, both partners (realistically) obtain the profits of the outside option from there on. Second, we assume that the per-period profits cannot be increased through cooperation in the product market. We do this to focus on the effects of repeated interaction within the alliance and not in the product market. In essence, we assume that firms do not coordinate product market decisions.

Success or failure We now specify which of the possible alliance outcomes, i.e. set of decisions taken by both partners for the entire game, constitute a success or a failure; these characterizations will be, among other performance indicators, also used to assess the accomplishment of alliances in our experimental results. We consider an alliance outcome to be "successful" if both partners cannot obtain higher payoffs in another alliance outcome. Consequently, we define a particular alliance-outcome as a "failure" if both partners could have done better. Formally,

Definition 1. Success and failure
An alliance outcome is a *success* if it is not Pareto-dominated by any other outcome, and a *failure* if it is Pareto-dominated.

Figure 2 displays the average per-period payoff that each partner can obtain in CD- and CP-alliances. In CD-alliances, the only alliance-outcome that is not Pareto-dominated is the one consisting in both partners participating and contributing resources in all periods (with an average payoff of 5). In CP-alliances, any alliance outcome that results in an average payoff in the two most outward lines would be considered a success. Among those, the only symmetric one (in which both players obtain the same payoff, as displayed in the dashed line), involves both partners participating and contributing resources in all periods. We will refer to the outcome where both partners participate and contribute resources in a given period (outcome $[C, C]$) as a *collaborative* outcome. The only symmetric successful alliance-outcome is, thus, the outcome where partners *collaborate* in every period. As shown in the following corollary, with the assumed discount factor (0.9), this alliance outcome can be achieved in equilibrium in both the CD- and CP-alliances.

[Figure 2 about here.]

Proposition 3. Success and failure in alliances
In both the CD- and CP-alliance, the only symmetric successful outcome, which involves both partners participating and contributing resources in all periods, can be supported as a Subgame Perfect Nash equilibrium (SPNE).

The reasoning is similar to the one of the Folk theorem in infinitely repeated games. Indeed, in these games, and provided that the discount factor is large enough, any alliance outcome which yields, on average, at least the minmax payoff of the static setup can be sustained by SPNE strategies in the dynamic game, where the minmax payoff is the outcome that minimizes the maximum harm that can be inflicted by others (Fudenberg and Maskin 1986). Given the payoff structure of our framework, it is obvious that for all alliances this is the exit option (payoff of 3). The set of average per-period payoffs achieving more than the minmax payoff are depicted by the thick lines in Figure 2. The proof in Appendix B checks that, indeed, the discount factor we use (0.9) is large enough to sustain as an equilibrium both partners participating and contributing resources in all periods.

However, the minmax criterion allows for many more outcomes to be sustainable in equilibrium. For instance, it can also be an equilibrium if, for a number of periods, partners are alternating in some sequence between contributing and not contributing resources, as long as this yields better results on average than staying out. It can also be an equilibrium if partners participate for some periods in the alliance (and earn on average at least the outside option), but then subsequently exit. Finally, partners may also not start the alliance in the first place, as this yields exactly the minmax payoff. These multiple possible equilibrium outcomes, however, should come as no surprise, since in frequent interactions trust and understanding can remedy (or worsen) in principle issues of cooperation and coordination (Gulati, Wohlgezogen, and Zhelyazkov 2012).

4 Results

As discussed, we are especially interested in how firms start off their collaboration, how they continue, and how they react to a setback. Accordingly, we now present the results from our experimental analysis following this logic.

4.1 Implementation, Success and Failure

Implementation We ran six experimental sessions, three of which involved several CD-alliances and another three involved several CP-alliances. That is, in every session subjects played only one type of alliance, but participated in several of these. We identify sessions with a number (1, 2, or 3) followed by either CD or CP. In total, we have 146 CD-alliance outcomes and, because of slightly less participants in one of the sessions, 140 CP-alliance outcomes. Dates, number of participants, and other details explained below are summarized in Table 4. All sessions were run at the experimental economics laboratory (LeeX) of Universitat Pompeu Fabra, in Barcelona, with subjects being students of all disciplines at this university. Pilot sessions were run at Bocconi University, in Milan. All sessions were run in English. Instructions, and choices were all electronically administered using the software z-tree (Fischbacher 2007).

[Table 4 about here.]

Table 4 gives some basic information about our experiments. From the basic descriptives, we can see that the average payoff of the CP-alliances was lower. Table 5 gives further descriptives on the fraction of symmetric successes ($[C, C]$ always) and terminated alliances (E played at some point). A similar pattern emerges: CP-alliances have a lower fraction of successes and a higher fraction of termination at some point in the experiment (these differences are significant as shown by the third column).[14] This higher fraction of termination leads CP-alliances to have a shorter duration. This pattern can also be seen in figure 3; partners in CP-alliances leave in general in earlier periods. A simple regression confirms that CP-alliances last significantly less time than CD-alliances (Table 6).

[Table 5 about here.]

[Figure 3 about here.]

[Table 6 about here.]

[14] Recall that in CP-alliances it is possible to have asymmetric success if one partner always contributes resources to the alliance while the other sometimes chooses to not contribute. Such successful alliances are thus characterized by outcome $[C, nC]$ (or $[nC, C]$, but not both!) in some periods and outcome $[C, C]$ in other periods. There are 13 CP-alliances with such an outcome in our experiment, and their consideration increases the fraction of successful CP-alliances to 0.414, still statistically smaller than the corresponding fraction for CD-alliances (Pearson's χ^2 with $p < 0.001$).

These statistics show that a CP-alliance seems less likely to generate successful stable alliances. However, it is interesting to study the dynamics of alliances in more detail, and in particular consider the pre-formation and alliance continuation stages. Hence, given the particular importance of the first period for future periods, we distinguish between the initial phase ($t = 1$) and all subsequent phases. This lets us capture some of the richness of the inter-temporal dynamics in alliances. Finally, we also look at the ability of alliances to revert to collaboration following a setback (alliance resilience).

4.2 First Period Outcomes

Looking at the left hand side bars of figure 4a, we can see that a higher share of partners does not participate in CP-alliances than in CD-alliances in the first period. Interestingly, over 6 percent of participants in CP-alliances choose exit already in the first period (whereas less than 2 percent of CD-alliances exits in the first period). These participants were not deterred by any past misbehavior on the part of their opponent, as this would have been their first encounter. Nor were they attracted by the opportunity not to contribute resources and thus get higher payoffs (i.e., playing nC while expecting the partner to choose C at $t = 1$). Indeed, exiting immediately appears sensible only if a partner expects the ongoing alliance to settle on non-collaboration immediately. The differences in exit between the coordination and cooperation alliances are statistically significant, as can be seen in table 7.[15]

[Figure 4 about here.]

Moreover, we see for those alliances that effectively start that nC is played by more than 30 percent of partners in CP-alliances, whereas this is less than 9 percent in CD-alliances; this difference is statistically significant (see table 7 and the left hand side bars of figure 4b). The lower exit and non-contribution (nC) frequencies of participants in CD-alliances translates in higher frequencies of choosing C. Indeed, almost 90 percent of first-period choices in CD-alliances are to contribute resources, whereas this number is slightly above 61 percent in CP-alliances.

[15] The overall Pearson's χ^2 test between first period exit choices is 8.24, with a p-value of 0.004. It must be noted that exit in the first period is, in general, an unattractive choice and remains so throughout the experiment: on average, in a single iteration there are 3.8 exits in the first period across both types of alliances (such small numbers lead to the paradoxical finding that in any given iteration there is no significant difference between alliance types, while across iterations the difference is significant).

[Table 7 about here.]

Overall, these observations reveal important differences in the initial conditions for co-operation and coordination alliances; CD-alliances start off much better. That is, partners in CD-alliances overwhelmingly start by playing C recognizing that there is value in a solution where both partners contribute relevant resources, and that this solution is attainable. Conversely, the free-riding incentive in a CP-alliance motivates almost one third of players in CP-alliances to play nC in the first round, or to exit immediately and thus preemptively leave the alliance to avoid losing from an unsuccessful alliance.

Given the comparably bad start CP-alliances have, what are their subsequent prospects? In other words, are initially unsuccessful alliances able to turn the cards around? Further, what happens with alliances that do get off on a good start? Are partners able to maintain collaborative efforts, or do they drift away to less good outcomes and eventually a premature exit? Choices in subsequent periods indicate that differences across alliances stay over time, as the right hand bars of figures 4a and 4b tentatively indicate. We explore this in more detail in the next subsection.

4.3 Continuation phase: the Dynamics of Successful Alliances

In a first step, we investigate from which period onwards successful alliances start collaborating. As can be seen from figure 5, alliances that manage to collaborate until the end overwhelmingly do so from the initial stage.

[Figure 5 about here.]

This is confirmed when looking at whether initially collaborative alliances can maintain this collaboration. If we look at what fraction of alliances that attain outcome $[C, C]$ in the first period still attain it later on, both types of alliances are good in maintain the good outcome where both partners contribute resources (Figure 6).

[Figure 6 about here.]

In a next step, to determine wether initial bad conditions can be turned around we look at the *percentage of alliances who do not collaborate in period 1* but then actually proceed to collaborating at some later point and from then on always collaborate. This percentage is 13.8 percent in CD-alliances and 12.8 percent in CP-alliances (not significantly different).

Therefore, partners in both types of alliances are not very good in turning bad initial situations around.

In sum, we have seen that the initial period determines the long-term success of an alliance to a significant extent for both types of alliances. We now look more in detail at the flipside, i.e. the dynamics of alliances that face a (temporal or permanent) breakdown in collaboration, in the next subsection.

4.4 Resilience phase: the Dynamics of Unsuccessful Alliances

Despite the generally lower level of collaboration in CP-alliances, many of the properties are similar across alliance types in the sense that initially good conditions are key for overall success. They do, however, differ significantly in their resilience, i.e. their reactions to one of the two partners not contributing resources (i.e., choosing nC). We distinguish between subsequent behavior to outcomes $[C, nC]$, $[nC, C]$ or $[nC, nC]$. These are the outcomes where an alliance suffers from a lack of contribution of resources from one partner, or both. Thereafter, we analyze the events leading up to a dissolution of an alliance, i.e. to partners choosing E.

In figure 7, the two off-diagonal graphs indicate that subjects in the two alliance types do behave differently, especially for $[C, nC]$. This is an outcome where the subject whose behavior is being studied chooses C while his/her partner chooses nC in the previous period. We can see that after an outcome of $[C, nC]$, a subject chooses C 40 percent of the time in a CD-alliance, but only 27 percent of the time in CP-alliances.[16] Moreover, after the outcome $[C, nC]$, a subject chooses nC 44 percent of the time in CD-alliances and 61 percent of the time in CP-alliances.[17] Thus, partners in CD-alliances are significantly more forgiving than partners in CP-alliances when their partner does not collaborate.

[Figure 7 about here.]

Conversely, the actions chosen after outcomes $[nC, C]$ and $[nC, nC]$ are not statistically different for both types of alliances. Thus, different dynamics for coordination and cooperation alliances are present after non-collaboration but are solely attributable to one partner being more forgiving after having observed non-collaboration from the other.

[16] The difference is significant at the 10% level, using both the two-sided Pearson's χ^2 test (2.86, p-value of 0.091) or the one-sided Fisher exact test ($p = 0.066$).

[17] The difference is significant at 5% for both the two-sided Pearson's χ^2 (4.32, $p = 0.038$) and the one-sided Fisher exact test ($p = 0.028$).

While there is no significant difference across alliance types of the frequency of E after any given outcome, one difference emerges when one considers longer histories (figure 8). While only *one* CD-alliance terminates the relationship in the second period, almost one third of all exit choices are made in the second period for CP-alliances, following an outcome where either partner chose nC.

This reinforces what we already found when looking at first period choices: in CP-alliances, partners anticipate further poor performance – especially after one instance of nC – and thus exit. As can be seen in figure 8, when a subject chooses to exit after outcome $[C, nC]$, in CP-alliances this is significantly more likely to be the first occurrence of nC.[18] In other words, subjects in cooperation alliances are less patient (or forgiving) and exit sooner after their partner has cheated on them.

[Figure 8 about here.]

The above findings are complemented by the observation that in CP-alliances it is more likely that the partner who first chooses to 'cheat' (first chooses nC) also is the partner who exits the alliance (first chooses E). Indeed, 42 percent of partners who first choose to not contribute resources also first choose exit in CP-alliances, while this is only 25 percent in CD-alliances.

In sum, partners in CD-alliances are (i) more forgiving towards cheating of their partner and, (ii) less inclined to cheat and subsequently exit an alliance. Therefore – although most of the differences across alliance types come from initial period choices – CD-alliances are also better armed to deal with momentary breakdowns in collaboration than CP-alliances.

5 Discussion and Conclusion

We show that alliances can face different problems, depending on the relationship between alliance partners ex-ante. If alliance partners are close competitors in the product market, alliances are more likely to face incentives problems, i.e. each player has an incentive to hold back effort in every stage of the alliance to free ride on its rival's efforts. Conversely, if alliance partners face large complementarities in their efforts, the problem is likely to be one of coordination, i.e. both alliance partners would contribute, but only if the partner

[18] The distribution of outcomes two periods before exit is chosen, including the possibility that *nothing* occurred (i.e., exit happened in the second period), conditional on outcome $[C, nC]$ one period before exit is chosen, is significantly different across alliance types as given by Fisher's exact test ($p = 0.002$).

also does so. We run a series of experiments on these two alliance types to compare their performance along three dimensions: their initial level of collaboration, their dynamic behavior, and their ability to bounce back from a breakdown in collaboration. We find that although coordination alliances do better in all three dimensions, the key difference lies in the way the alliance begins. In cooperation alliances, a significant fraction of alliances start with an attempt to cheat by one or both of the partners, making it difficult to recover. Conversely, in coordination alliances the level of collaboration is significantly higher at the beginning. Even among the alliances that start out collaborating, coordination alliances are more likely to continue collaborating, although the difference in alliance continuation is less marked than the difference in initial collaboration. Finally, both alliance types do not bounce back easily from a breakdown in collaboration. Again however, coordination alliances perform somewhat better than cooperation ones.

These results have several implications. First, the early stages of an alliance matter most in determining its long-term success. At first, this may not sound surprising, but given the sizable literature on firms' ability to learn to ally, our results suggest that starting cautiously and hoping for increased collaboration over time may not lead to success in the long-run. This is because alliances with a prior negative history will likely continue to fail. Second, alliance partners should design safeguards against a breakdown in collaboration, as it is hard to come back from a setback. Third, short of actively working to change the degree of complementarities and the degree of product market competition, which may simply be impossible in some settings, choosing the right alliance partner matters. Specifically, a partner who shares enough commonalities to generate complementarities while being distant enough so competition is not too fierce would likely resemble a coordination (CD-)alliance, which performed better in our experiments.

Our paper only starts to capture some of the rich dynamics of real-world strategic alliances. By using both the conceptual clarity of a formal model, however, and the behavioral realism of an experiment, we believe that we can add valuable insights both into the microfoundations of strategic alliances and in their subsequent dynamics. Finally, by presenting some initial results on the ability of alliances to revert to collaboration after a setback, we hope to stimulate further research on the resilience of alliance types as another relevant dimension of alliance performance.

References

Agarwal, Rajshree, Rachel Croson, and Joseph Mahoney (2010). “The role of incentives and communication in strategic alliances: an experimental investigation”. In: *Strategic Management Journal* 31, pp. 413–437.

Agarwal, Rajshree et al. (2012). “Spillovers Across Organizational Architectures: The Role of Prior Resource Allocation and Communication in Post-Acquisition Coordination Outcomes”. In: *Strategic Management Journal* 33, pp. 710–733.

Amaldoss, Wilfred and Richard Staelin (2010). “Cross-function and same-function alliances: how does alliance structure affect the behavior of partnering firms?” In: *Management Science* 56.2, pp. 302–317.

Anand, Bharat and Tarun Khanna (2000). “Do Firms Learn to Create Value? The Case of Alliances”. In: *Strategic Management Journal* 21.3, pp. 295–315.

Arend, Richard (2009). “Reputation for cooperation: contingent benefits in alliance activity”. In: *Strategic Management Journal* 30, pp. 371–385.

Arend, Richard and Darryl Seale (2005). “Modeling Alliance Activity: An Iterated Prisoners’ Dilemma with Exit Option”. In: *Strategic Management Journal* 26, pp. 1057–1074.

Cachon, Gerard P. and Colin F. Camerer (1996). “Loss-Avoidance and Forward Induction in Experimental Coordination Games”. In: *The Quarterly Journal of Economics* 111.1, pp. 165–194.

Crocker, Keith and Kenneth Reynolds (1993). “The Efficiency of Incomplete Contracts: An Empirical Analysis of Air Force Engine Procurement”. In: *The RAND Journal of Economics* 24.1, pp. 126–146.

Deeds, David and Charles Hill (1996). “Strategic Alliances and the Rate of New Product Development: an Empirical Study of Entrepreneurial Biotechnology Firms”. In: *Journal of Business Venturing* 11.1, pp. 41–55.

Doz, Yves (1996). “The Evolution of Cooperation in Strategic Alliances: Initial Conditions or Learning Processes?” In: *Strategic Management Journal* 17, pp. 55–83.

Doz, Yves and Gary Hamel (1998). *Alliance Advantage: The Art of Creating Value Through Partnering*. Harvard Business Press.

Dyer, Jeffrey H. (1996). “Does Governance Matter? Keiretsu Alliances and Asset Specificity as Sources of Japanese Competitive Advantage”. In: *Organization Science* 7.6, pp. 649–666.

Dyer, Jeffrey H. and Harbir Singh (1998). “The Relational View: Cooperative Strategy and Sources of Interorganizational Competitive Advantage”. In: *Academy of Management Review* 23.4, pp. 660–679.

Fischbacher, Urs (2007). “z-Tree: Zurich Toolbox for Ready-made Economic Experiments”. In: *Experimental Economics* 10, pp. 171–178.

Fudenberg, Drew and Eric Maskin (1986). “The Folk Theorem in Repeated Games with Discounting or with Incomplete Information”. In: *Econometrica* 54.3, pp. 533–554.

Fudenberg, Drew and Jean Tirole (1984). “The Fat Cat Effect, The Puppy-Dog Ploy, and The Lean and Hungry Look”. In: *American Economic Review* 74.2, pp. 361–366.

Garrette, Bernard, Xavier Castañer, and Pierre Dussauge (2009). “Horizontal Alliances as an Alternative to Autonomous Production: Product Expansion Mode in the Worldwide Aircraft Industry 1945-2000”. In: *Strategic Management Journal* 30.885-894.

Gibbons, Robert (1992). *Game Theory for Applied Economists*. Princeton University Press.

Gilbert, Richard (2006). “Looking for Mr. Schumpeter: Where are we in the competition-innovation debate?” In: *Innovation Policy and the Economy*. Ed. by Adam Jaffe, Josh Lerner, and Scott Stern. Vol. 6. The MIT Press.

Gottschlag, Oliver and Maurizio Zollo (2007). “Interest Alignment and Competitive Advantage”. In: *Academy of Management Review* 32.2, pp. 418–437.

Gulati, Ranjay (1995). “Does Familiarity Breed Trust? The Implications of Repeated Ties for Contractual Choice in Alliances”. In: *Academy of Management Journal* 38.1, pp. 85–112.

Gulati, Ranjay, Paul Lawrence, and Phanish Puranam (2005). “Adaptation in Vertical Relationships: Beyond Incentive Conflict”. In: *Strategic Management Journal* 26.5, pp. 415–440.

Gulati, Ranjay and Harbir Singh (1998). “The Architecture of Cooperation: Managing Coordination Costs and Appropriation Concerns in Strategic Alliances”. In: *Administrative Science Quarterly* 43.4, pp. 781–814.

Gulati, Ranjay and Maxim Sytch (2007). “Dependence Asymmetry and Joint Dependence in Interorganizational Relationships: Effects of Embeddedness On a Manufacturer’s

Performance in Procurement Relationships". In: *Administrative Science Quarterly* 52.1, pp. 32–69.

Gulati, Ranjay, Franz Wohlgezogen, and Pavel Zhelyazkov (2012). "The Two Facets of Collaboration: Cooperation and Coordination in Strategic Alliances". In: *Academy of Management Annals* 6.1, pp. 531–583.

Hagedoorn, John (1993). "Understanding the Rationale of Strategic Technology Partnering: Interorganizational Modes of Cooperation and Sectoral Differences". In: *Strategic Management Journal* 14.5, pp. 371–385.

Hahn, Jinyong, Petra Todd, and Wilbert Van der Klaauw (2001). "Identification and Estimation of Treatment Effects With a Regression-Discontinuity Design". In: *Econometrica* 69.1, pp. 201–209.

Hamel, Gary (1991). "Competition For Competence and Interpartner Learning Within International Strategic Alliances". In: *Strategic Management Journal* 12.1, pp. 83–103.

Harrigan, Kathryn Rudie (1985). "Vertical Integration and Corporate Strategy". In: *Academy of Management Journal* 28.2, pp. 397–425.

Hill, Robert and Don Hellriegel (1994). "Critical Contingencies in Joint Venture Management: Some Lessons From Managers". In: *Organization Science* 5.4, pp. 594–607.

Kale, Prashant, Jeffrey H. Dyer, and Harbir Singh (2002). "Alliance Capability, Stock Market Response, and Long-Term Alliance Success: The Role of the Alliance Function". In: *Strategic Management Journal* 23.8, pp. 747–767.

Kale, Prashant and Harbir Singh (2009). "Managing Strategic Alliances: What Do We Know Now, and Where Do We Go From Here?" In: *Academy of Management Perspectives* 23.3, pp. 45–62.

Kale, Prashant, Harbir Singh, and Howard Perlmutter (2000). "Learning and Protection of Proprietary Assets in Strategic Alliances: Building Relational Capital". In: *Strategic Management Journal* 21.3, pp. 217–237.

Khanna, Tarun, Ranjay Gulati, and Nitin Nohria (1998). "The Dynamics of Learning Alliances: Competition, Cooperation, and Relative Scope". In: *Strategic Management Journal* 19.3, pp. 193–210.

Kogut, Bruce (1988). "Joint Ventures: Theoretical and Empirical Perspectives". In: *Strategic Management Journal* 9.4, pp. 319–332.

Kogut, Bruce (1989). "The Stability of Joint Ventures: Reciprocity and Competitive Rivalry". In: *Journal of Industrial Economics* 38.2, pp. 183–198.

Lunnan, Randi and Sven Haugland (2008). "Predicting and Measuring Alliance Performance: A Multidimensional Analysis". In: *Strategic Management Journal* 29.5, pp. 545–556.

Mesquita, Luiz, Jaideep Anand, and Thomas Brush (2008). "Comparing the Resource-Based and Relational Views: Knowledge Transfer and Spillover in Vertical Alliances". In: *Strategic Management Journal* 29.9, pp. 913–941.

Oxley, Joanne E. (1997). "Appropriability Hazards and Governance in Strategic Alliances: A Transaction Cost Approach". In: *Journal of Law, Economics, and Organization* 13.2, pp. 387–409.

Park, Seung Ho and Michael Russo (1996). "When Competition Eclipses Cooperation: An Event-history Analysis of Joint Venture Failure". In: *management Science* 42.6, pp. 875–890.

Park, Seung Ho and Gerardo Ungson (1997). "The Effect of National Culture, Organizational Complementarity, and Economic Motivation on Joint Venture Dissolution". In: *Academy of Management Journal* 40.2, pp. 279–307.

— (2001). "Interfirm Rivalry and Managerial Complexity: A Conceptual Framework of Alliance Failure". In: *Organization Science* 12.1, pp. 37–53.

Parkhe, Arvind (1993). "Strategic Alliance Structuring: a Game Theoretic and Transaction Cost Examination of Interfirm Cooperation". In: *Academy of Management Journal* 36.4, pp. 749–829.

Porter, Michael and Mark Fuller (1986). "Coallitions and Global Strategies". In: *Competition in Global Industries*. Ed. by Michael Porter. Boston: Harvard Business School Press, pp. 315–343.

Rafique, Hashmi Aamir (2013). "Competition and Innovation: The Inverted-U Relationship Revisited". In: *Review of Economics and Statistics* 95.5, pp. 1653–1668.

Reuer, Jeffrey J. and Maurizio Zollo (2005). "Termination Outcomes of Research Alliances". In: *Research Policy* 34, pp. 101–115.

Robinson, David and Toby Stuart (2007). "Financial Contracting in Biotech Strategic Alliances". In: *Journal of Law and Economics* 50.3, pp. 559–596.

Sampson, Rachelle (2004). “Organizational Choice in R&D Alliances: Knowledge-Based and Transaction Cost Perspectives”. In: *Managerial and Decision Economics* 25.6-7, pp. 421–436.

Silverman, Brian S. and Joel Baum (2002). “Alliance-Based Competitive Dynamics”. In: *Academy of Management Journal* 45, pp. 791–806.

Teece, David (1988). “Technological change and the nature of the firm”. In: *Technical change and economic theory*. Ed. by Giovanni Dosi et al. London: Pinter Publishers.

Uzzi, Brian (1997). “Social Structure and Competition in Interfirm Networks: The Paradox of Embeddedness”. In: *Administrative Science Quarterly* 42.1, pp. 35–67.

Varshney, Lav and Daniel Oppenheim (2011). “On cross-enterprise collaboration”. In: *Business Process Management*. Berlin: Springer, pp. 29–37.

Vives, Xavier (2008). “Innovation and Competitive Pressure”. In: *Journal of Industrial Economics* 56.3, pp. 419–469.

Wassmer, Ulrich and Pierre Dussauge (2011). “Value creation in alliance portfolios: The benefits and costs of network resource interdependencies”. In: *European Management Review* 8.1, pp. 47–64.

Williamson, Oliver (1985). *The Economic Institutions of Capitalism*. New York: Free Press.

Zaheer, Akbar, Bill McEvily, and Vincenzo Perrone (1998). “Does Trust Matter? Exploring the Effects of Interorganizational and Interpersonal Trust on Performance”. In: *Organization Science* 9.2, pp. 141–159.

Zeng, Ming and Xiao-Ping Chen (2003). “Achieving cooperation in multiparty alliances: A social dilemma approach to partnership management”. In: *Academy of Management Review* 28.4, pp. 587–605.

Table 1: Profits as a function of the decision to contribute resources

		Firm j	
		$\boldsymbol{C}$	$\boldsymbol{nC}$
Firm i	$\boldsymbol{C}$	$\pi\,(c-kr,d)-e$ $\pi\,(c-kr,d)-e$	$\pi\,(c-r,d)$ $\pi\,(c-r,d)-e$
	$\boldsymbol{nC}$	$\pi\,(c-r,d)-e$ $\pi\,(c-r,d)$	$\pi\,(c,d)$ $\pi\,(c,d)$

Table 2: Profits as a function of participation and contribution decisions

	Firm j		
Firm i	$\boldsymbol{C}$	$\boldsymbol{nC}$	$\boldsymbol{E}$
$\boldsymbol{C}$	$\pi\,(c-kr,d)-e-F$ $\pi\,(c-kr,d)-e-F$	$\pi\,(c-r,d)-F$ $\pi\,(c-r,d)-e-F$	$\pi\,(c,d)$ $\pi\,(c,d)$
$\boldsymbol{nC}$	$\pi\,(c-r,d)-e-F$ $\pi\,(c-r,d)-F$	$\pi\,(c,d)-F$ $\pi\,(c,d)-F$	$\pi\,(c,d)$ $\pi\,(c,d)$
$\boldsymbol{E}$	$\pi\,(c,d)$ $\pi\,(c,d)$	$\pi\,(c,d)$ $\pi\,(c,d)$	$\pi\,(c,d)$ $\pi\,(c,d)$

Table 3: Experimental alliances: (a) $CD-$alliance. (b) $CP-$alliance

	Firm j		
Firm i	$\boldsymbol{C}$	$\boldsymbol{nC}$	$\boldsymbol{E}$
	5	4	3
$\boldsymbol{C}$	5	-1	3
	-1	2	3
$\boldsymbol{nC}$	4	2	3
	3	3	3
$\boldsymbol{E}$	3	3	3

(a)

	Firm j		
Firm i	$\boldsymbol{C}$	$\boldsymbol{nC}$	$\boldsymbol{E}$
	5	6	3
$\boldsymbol{C}$	5	-1	3
	-1	2	3
$\boldsymbol{nC}$	6	2	3
	3	3	3
$\boldsymbol{E}$	3	3	3

(b)

Table 4: Setup of the experimental alliances

Session	*Date*	*Alliance*	*# of subjects*	*# of alliances*	*Periods in each alliance*	*Mean and Variance of payoff*
$1CD$	27-09-10	CD	20	6	$(1, 12, 6, 9, 3, 4)$	(12.74, 0.82)
$1CP$	28-09-10	CP	20	–	–	(11.41, 0.80)
$2CD$	22-11-10	CD	18	4	$(13, 3, 5, 18)$	(12.96, 0.97)
$2CP$	30-11-10	CP	20	–	–	(12.50, 1.05)
$3CD$	18-03-11	CD	20	5	$(11, 5, 10, 9, 10)$	(13.63, 1.95)
$3CP$	18-03-11	CP	16	–	–	(11.64, 0.30)

Notes: CP alliances had the same number of periods as in the equally-numbered CD session. Mean and variance of subjects' take-home pay (in Euros) are presented as *(Mean, Variance)*; see the experiment instructions in Appendix C for the exact payment details.

Table 5: Fraction of successful and fraction of terminated CD- and CP-alliances

	CD-alliance	CP-alliance	Pearson's χ^2 test
Freq. of Successful Alliances	0.7192	0.3214	45.33 $p < 0.001$
Freq. of Terminated Alliances	0.2055	0.4857	24.92 $p < 0.001$

Notes: Sample consists of 146 observations in CD-alliances and 140 in CP-alliances.

Table 6: Duration of alliances

Intercept	CP-alliance
0.8321**	−0.1722**
(0.0485)	(0.0372)

Notes: Regression of percentage of total alliance periods before exit (*period of exit/period of random termination*) on alliance type, where *CD-alliance* is the base and *CP-alliance* is included as a categorical variable. Alliance controls are included but not reported. The sample has 286 observations. ** Significance at 1% level.

Table 7: Choices in the first period of an alliance

Choice:	CD-alliance	CP-alliance	Pearson's χ^2 test
C	0.8973	0.6179	61.26 $p < 0.001$
nC	0.0856	0.3178	48.31 $p < 0.001$
E	0.0171	0.0643	8.24 $p = 0.004$

Notes: Sample consists of 146 observations in CD-alliances and 140 in CP-alliances.

Figure 1: Alliance types

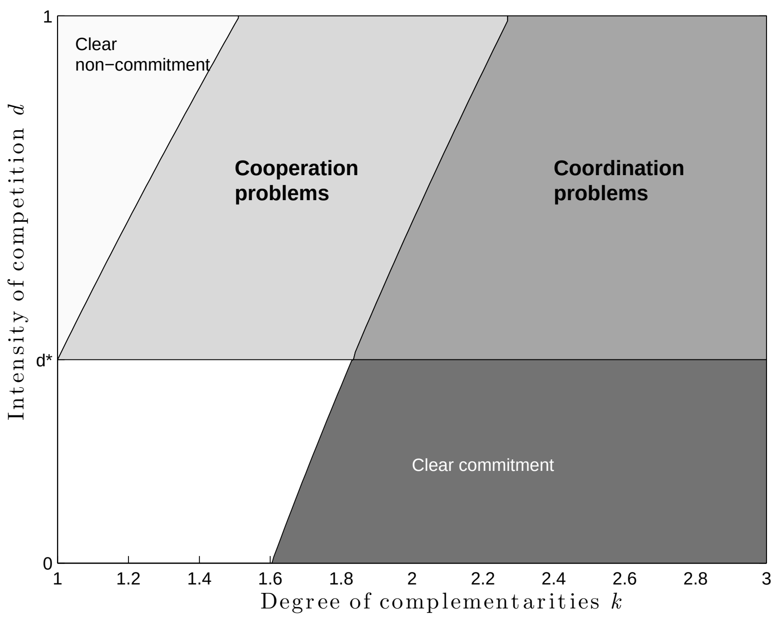

Notes: Types of strategic interaction and equilibrium resource contributions as a function of the degree of competition (d) and degree of complementarities (k) for the Cournot-quantity competition model. In terms of parameters, we take the markup $(a - c) = 1$, the cost reduction if only one firm contributes $r = 0.25$, and the cost of contributing resources to the alliance $e = 0.1$. See Appendix A for a full description of this model.

Figure 2: Set of average per-period payoffs that each partner can obtain in any CD- and CP-alliance outcome (in black and red, respectively). In bold, those that are above the minmax payoff (3,3)

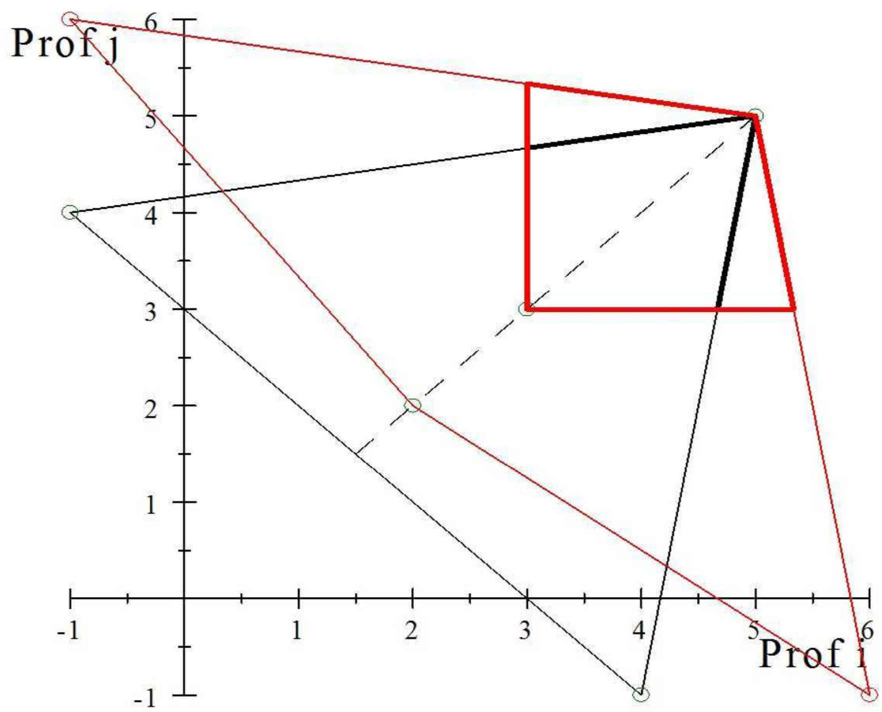

Figure 3: Fraction of all first choices of exit that happen in a given period

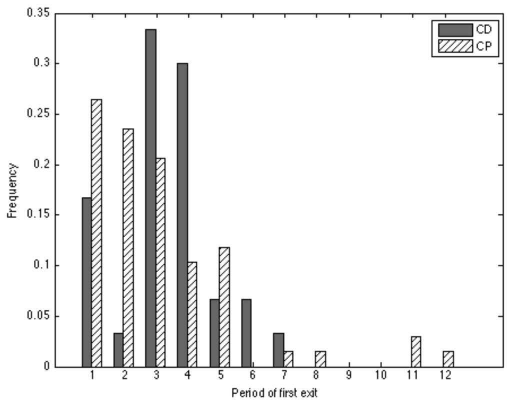

Figure 4: Per-period choices in CD- and CP-alliances.

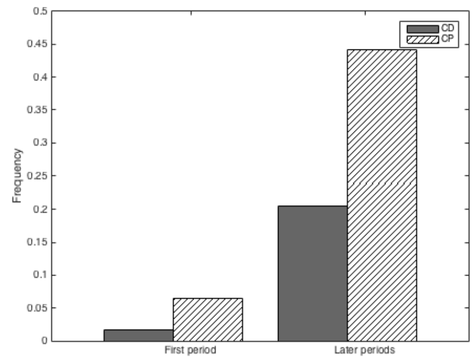

(a) Exit as a fraction of all per-period choices.

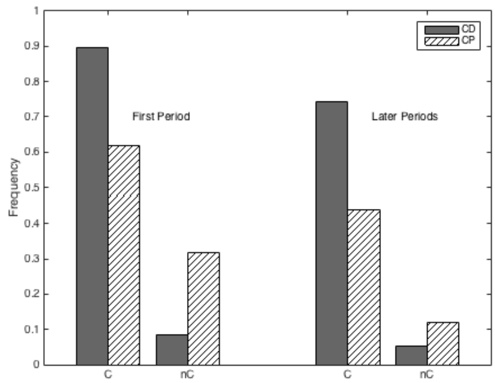

(b) Choices C and nC as a fraction of all per-period choices in ongoing alliances. *First Period* considers only choices in the first period of every alliance.

Figure 5: Among alliances who play $[C, C]$ until the last period of the interaction, what fraction started uninterrupted play of $[C, C]$ in a given period?

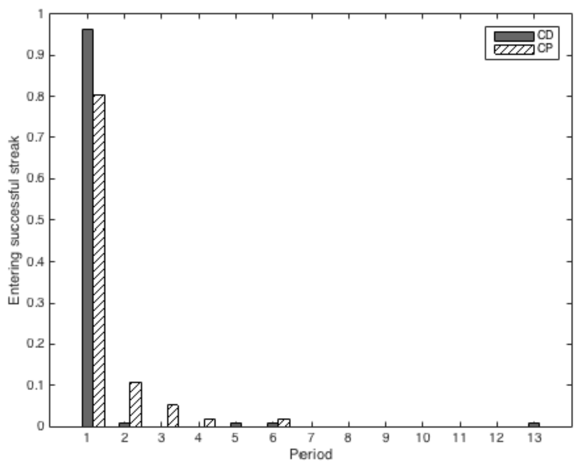

Figure 6: Percentage of all alliances that played $[C, C]$ in the first period and either: (i) still play $[C, C]$ in any given period or (ii) played $[C, C]$ until random termination of the alliance

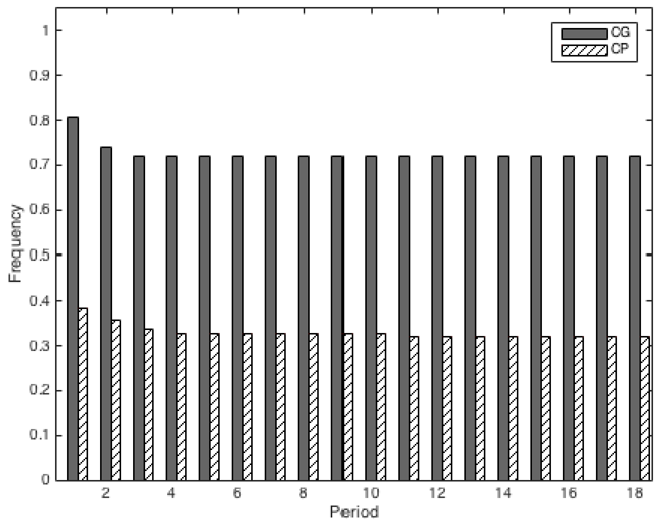

Figure 7: Frequency of play of each action in period t after each possible outcome in period $t-1$, by type of alliance game.

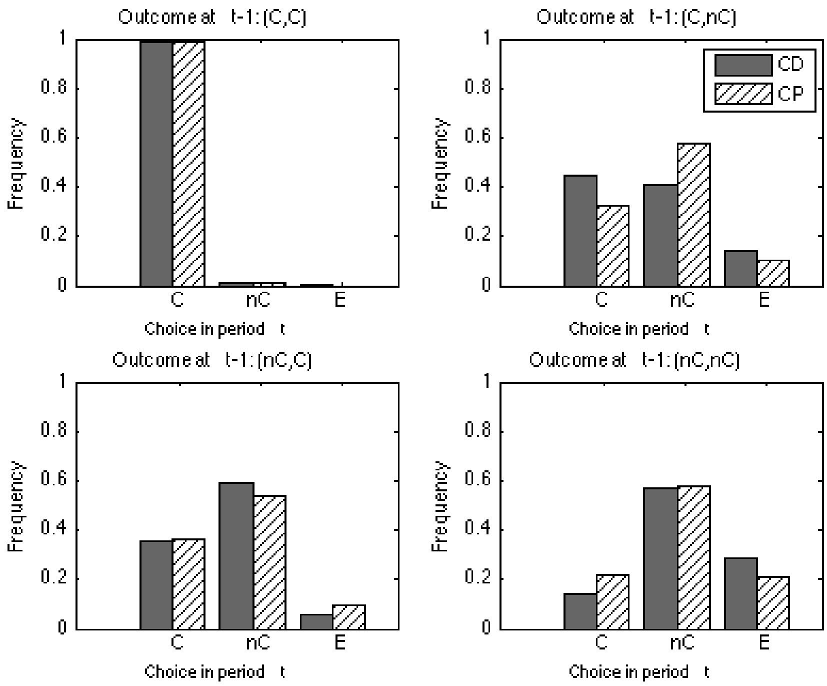

Notes: Outcomes involving a choice of E are automatically followed by E and are therefore not considered.

Figure 8: Joint histogram of outcomes one period and two periods prior to an exit that happens in any period later than the first. Outcome "N/A" two periods prior two exit represents alliances where exits occurs in the second period of the interaction.