# TESTING MODELS OF STRATEGIC UNCERTAINTY: EQUILIBRIUM SELECTION IN REPEATED GAMES

**Marta Boczoń**
Copenhagen Business School

**Emanuel Vespa**
UC San Diego

**Taylor Weidman**
University of Pittsburgh

**Alistair J. Wilson**
University of Pittsburgh

**Abstract**
In repeated games, where both collusive and noncollusive outcomes can be supported as equilibria, it is crucial to understand the likelihood of selection for each type of equilibrium. Controlled experiments have empirically validated a selection criterion for the two-player repeated prisoner's dilemma: the basin of attraction for always defect. This prediction device uses the game primitives to measure the set of beliefs for which an agent would prefer to unconditionally defect rather than attempt conditional cooperation. This belief measure reflects strategic uncertainty over others' actions, where the prediction is for noncooperative outcomes when the basin measure is full, and cooperative outcomes when empty. We expand this selection notion to multi-player social dilemmas and experimentally test the predictions, manipulating both the total number of players and the payoff tensions. Our results affirm the model as a tool for predicting long-term cooperation, while also speaking to some limitations when dealing with first-time encounters. (JEL: C73, C92, D91)

## 1. Introduction

Identifying which of many possible equilibria best captures economic behavior is of central importance for applications with repeated interactions. For example, in models of oligopoly, both collusive and noncollusive equilibria can arise. To help guide assumptions over equilibrium selection, experimental work has sought to uncover simple theoretical criteria that can predict the likelihood of collusion based on primitives such as payoffs and discount rates. Thus far, the main body of experimental work on selection has focused on the canonical two-player indefinitely repeated

prisoner's dilemma (RPD). However, it is unknown to what extent the predictive criteria from two-player environments can be used to predict selection in games with more than two players.

The *basin of attraction for always defect* (Blonski and Spagnolo, 2001, 2015) has been shown to simply organize experimental data in a meta-study of two-player RPD games (Dal Bó and Fréchette, 2018). The measure's calculation inputs are stage-game payoffs and the discount factor. The measure's output is the set of beliefs on the other player choosing to conditionally cooperate for which permanent defection is a best response. The wider the set of beliefs where defection is a best response, the greater the risk in attempting to cooperate, which is why this measure is thought of as a proxy for uncertainty over others' behavior (i.e., strategic uncertainty). Experimental data starting with Dal Bó and Fréchette (2011) show that when the theoretical size of the always defect basin is high (low), observed cooperation rates tend to be low (high). Furthermore, when the basin size is less than (greater than) half, it aligns with the concept of risk dominance. Therefore, this simple ordinal property serves as a clear line-in-the-sand for predicting regions where we expect/do not expect collusive outcomes.

Our paper focuses on a simple and relevant extension of the model to more than two players. In environment with $N$ players, an agent must assess the chances that multiple other players will cooperate. We develop two benchmarks to extend the measure of strategic uncertainty. A natural theoretical benchmark is an independent-belief extension, which considers symmetric and independent beliefs about each of the other players. At the other extreme, we also consider a setting where beliefs about other players are perfectly correlated. In the perfectly-correlated extension, the addition of another player does not impact strategic uncertainty, as the actions of other players in the game are assumed to be perfectly correlated. This extreme serves as a natural interpretation for a null hypothesis over the number of players, where the prediction for the $N$-player game will, ceteris paribus, be the same as for its two-player counterpart. With these two benchmarks, our experimental treatments allow us to isolate the effects of strategic uncertainty due to the higher $N$ relative to the standard two-player game.

Our experimental design provides directional and null predictions for each extension. To achieve this goal, we introduce a second treatment variable that the prior RPD literature highlights as a clear driver of behavior: the stage-game payoff gain that a player gets from a defection, $x$. For illustration, consider the effect of reducing $x$. A lower temptation payoff reduces strategic uncertainty according to both extensions, predicting higher cooperation. In particular, this second parameter provides for a directional prediction under the correlated extension. Further, by shifting both $N$ and $x$ together, we can generate null-effect treatment comparisons under our independent extension. That is, a predicted increase in strategic uncertainty from higher $N$ can be perfectly compensated by a reduction in $x$, generating a stronger test of this extension.

By manipulating $N$ and $x$, we create a $2 \times 2$ between-subjects design across our two basin extensions, generating directional and null predictions for each. This design enables us to assess which of the two extensions more effectively captures behavior. Furthermore, it remains a possibility that both extensions lack predictive power, indicating that coordination in a game involving more than two players may not be associated with strategic uncertainty.

Our core results indicate that the independent-basin extension best organizes the longer-run (i.e., *ongoing*) behavior. Under the independent-belief extension we observe large behavioral shifts in the predicted direction when varying $N$ in isolation. When manipulating $x$ and $N$ in opposite directions we observe no significant changes in behavior, which is in line with the prediction of the independent extension. However, while the independent extension succeeds in predicting the longer-run cooperation that may be the most relevant for applications, the measure is not a good predictor of intentions to cooperate, as captured by *initial* cooperation.

Our core findings suggest that equilibrium selection is driven by strategic uncertainty over the behavior of other players. Therefore, eliminating or minimizing doubts about others' actions should render the predictions from the basin model irrelevant. We explore this hypothesis in an additional treatment with pre-play communication. Here, participants have the opportunity to exchange free-form messages before the repeated game begins, a feature designed to reduce uncertainty about the strategic intentions of others (see Kartal and Müller, 2022). In a parameterization where initial cooperation rates are below one percent in the treatment without communication, the introduction of pre-play chat shifts behavior to the other extreme, resulting in initial (ongoing) cooperation rates of 95 (80) percent.

While we observe that predictions based on strategic uncertainty lose validity with explicit collusion, the model performs well in several robustness treatments with tacit collusion. We first assess the extent to which selected equilibria are sticky when game parameters change. Specifically, we introduce a group-size change halfway through an experimental session, transitioning the same participants from a four-player game to a two-player game, and vice versa. If the selected equilibrium in the first half is sticky, then varying $N$ will not affect behavior. As a consequence, the independent-basin extension, which better organizes results in our between-subjects comparison, would be irrelevant for comparative statics within a market. However, if strategic uncertainty changes with new parameters, then an increase (decrease) in $N$ should decrease (increase) beliefs in others' cooperation after the parameter change, altering cooperation. Our findings indicate some stickiness in behavior in the short run, but we do not observe stickiness in longer-term behavior. Cooperation levels adjust after a change in $N$, moving with experience toward the levels observed in the sessions with fixed $N$. These results validate the independent-extension measure as a predictor for equilibrium selection even within a particular context.

In our second robustness treatment, we relax the condition for group success. Our previous treatments require joint cooperation from all $N$ players for a group-wide success. Hence, increasing the number of players from two to four makes it mechanically harder to achieve success at any fixed rate of individual cooperation (i.e., two-from-two is easier than four-from-four). In this robustness treatment, however, we only require half of the players to cooperate for a group-wide success. Consider comparing the baseline reference (two-from two) to this robustness (two-from-four) treatment. The comparison increases $N$ from two to four, but this change does not make it mechanically *harder* to achieve success. That is, achieving a success in the two-from-two game is harder than in the two-from-four case. Despite easing the conditions for success in the robustness treatment, strategic uncertainty increases, and the independent-extension predicts lower cooperation. The reason behind this shift is that coordination in the two-from-four treatment becomes more challenging, as it introduces the question of which players must cooperate and which can free-ride. Consistent with the counter-intuitive prediction, our experimental results indicate that cooperation rates are lower under the two-from-four requirement for an efficient outcome than the standard RPD where we require two-from-two.

### 1.1. Literature

This paper is connected to several strands of the literature. Our design is based on the recent consolidation of the experimental RPD literature presented in Dal Bó and Fréchette (2018). While one of our baseline treatments replicates a standard finding in the literature,[1] we generalize the equilibrium selection model by adding an additional source of strategic uncertainty: the number of players, $N$.[2] Where the literature has developed this model for explanatory purposes, our approach is both to expand the model to a new setting, but also to test it as the core experimental object.

Our generalization of the strategic uncertainty model is carried out in two ways. The first extension (and most standard, given its use of independent beliefs) formalizes a distinct source of strategic uncertainty from the payoff-based source in the meta-study. An alternative extension (based on fully correlated beliefs) reflects a null effect, that the newly introduced source has no effect. As such, our generalization offers a potentially profitable design approach for future research examining other channels for

---

1.  As highlighted by Berry, Coffman, Hanley, Gihleb, and Wilson (2017), experimental replications can seem less frequent than they are if papers fail to advertise the features that are replications.

2.  The basin measure, detailed in Section 2, seeks to capture the intuition from Harsanyi and Selten (1988)'s risk dominance, and was initially proposed by Blonski and Spagnolo (2001, 2015). The basin measure was first empirically tested by Dal Bó and Fréchette (2011). See also Fudenberg, Rand, and Dreber (2012) for an examination of the effects with imperfect monitoring, Kartal and Müller (2022) for a test of a selection theory based on individual heterogeneity in preferences over dynamic strategies, and Mermer, Mueller, and Suetens (2021) for two-player games of strategic complements and substitutes.

strategic uncertainty effects—asymmetries in the action space or payoffs, the effects of sequentiality, etc.[3]

Our environment also allows us to better distinguish between empirical measures linked to the selection model. That is, using literature-level data assembled by Dal Bó and Fréchette (2018), we show that the two-player RPD strategic uncertainty model is suitable to predict both initial and ongoing cooperation.[4] However, with more than two players, this is no longer the case. Here, we demonstrate that the strategic uncertainty model is better suited to predict ongoing collusion rather than initial intentions to collude.[5]

This paper is part of a broader literature that seeks to understand and document regularities in equilibrium selection, in particular, regularities that are amenable to theoretical modeling. To this end, our measures of strategic uncertainty are particularly promising, as the equilibrium objects required for calculation are computationally simple: the stationary noncollusive equilibrium and the history-dependent collusive equilibrium. In environments beyond the RPD in which the equilibrium outcomes are held constant, the model can be similarly extended per our illustration with a move to $N$ players. However, in more complex environments with changing sets of equilibria, the constraint to two focal equilibria may lose validity and/or raise questions as to which two strategies are focal. Examples of more-complex settings include dynamic games in which the stage environment changes across supergames, and the space of strategies grows exponentially. Vespa and Wilson (2020) focus on a horse-race examination of which two equilibria are focal (from a wider set of possible alternatives) to rationalize behavior in dynamic games. That paper identifies a similar strategic uncertainty measure constructed around the most-efficient Markov perfect equilibrium and the best *symmetric* collusive equilibrium. A strategic-uncertainty model based on these strategies predicts behavior, where these strategies dovetail with repeated game strategies in the simpler environment studied here.[6]

---

3.    See Ghidoni and Suetens (2022) and Kartal and Müller (2022) for experimental examinations of the effect of sequentiality in RPD settings through a reduction in strategic uncertainty.

4.    With two players, the introduction of sequential moves adds extra variability for identification. Ghidoni and Suetens (2022) also find that ongoing measures are better predicted than initial rates.

5.    Ongoing cooperation is a measure that is likely to be more relevant for empirical applications where collusion may be a worry. For instance, from Harrington, Gonzalez, and Kujal (2016), page 256: "(...) collusion is more than high prices, it is a mutual understanding among firms to coordinate their behavior. (...) Firms may periodically raise price in order to attempt to coordinate a move to a collusive equilibrium, but never succeed in doing so; high average prices are then the product of failed attempts to collude."

6.    The applications of dynamic games are extensive, thanks to their inherent flexibility. The ongoing research on equilibrium selection in dynamic games builds upon recent work, among others, by Battaglini, Nunnari, and Palfrey (2012, 2016); Agranov, Fréchette, Palfrey, and Vespa (2016); Kloosterman (2019); Vespa and Wilson (2019); Rosokha and Wei (2020); Salz and Vespa (2020); Vespa (2020).

An experimental literature on behavior in oligopolies documents that collusion responds to the number of players. Both Cournot (Huck, Normann, and Oechssler, 2004; Horstmann, Krämer, and Schnurr, 2018) and Bertrand settings (Dufwenberg and Gneezy, 2000) indicate that as the number of players increases collusion becomes less likely, often as soon as $N$ exceeds two.[7] We contribute to this literature on two margins. First, we examine how changes to $N$ affect outcomes in an infinite horizon with collusive and noncollusive equilibria. Second, and crucially, we focus not only on the qualitative directional effects of $N$, but also, on validating the model suitability for studying strategic uncertainty. Specifically, the model, if validated, will help us understand the extent of substitutability between game primitives, which, in turn, can prove useful in predicting the directional effects of more-nuanced, multi-dimensional counterfactuals.

Our work is also related to the experimental literature on mergers that manipulates the number of players. As surveyed by Goette and Schmutzler (2009), some experiments deal with "pseudo-mergers," where a subset of the original firms remains in the market (see, for example, Huck, Konrad, Müller, and Normann, 2007). Other experiments implement "real mergers," where mergers introduce other changes in the market beyond $N$ (Davis, 2002). Our strategic-uncertainty measure can predict counterfactual behavior in both settings. Another discussion in this literature is whether merger effects are evaluated within the same group of participants (within-subject designs) or across different groups (between-subject designs). In this paper, we also conduct within-subject sessions at the same parameterization, demonstrating that although there can be meaningful short-run differences, with enough experience the results align.[8]

The effects of communication devices as a support for collusion are well established in the experimental literature. As surveyed in Cason (2008) and Harrington, Gonzalez, and Kujal (2016), more-structured, limited forms of communication usually result in small, temporary collusive gains, where free-form communication generates large, long-lasting effects.[9] For these reasons, we also examine unrestricted chat messages as a strong coordination device. Our collusive results indicate that the domain for our strategic-uncertainty measure based on tacit collusion does not include environments where explicit collusion is allowed. However, we show that there are clear limits on

---

7.   See also references in Potters and Suetens (2013) for similar findings.

8.   Differences in behavior tend to be stickier when changes are small or introduced gradually. Weber (2006) shows that gradually increasing the number of players in a coordination game yields different results relative to situations where the game begins with a large group. The gradual introduction of changes to the payoff primitives has also been shown to have effects in repeated games; see Kartal, Müller, and Tremewan (2021). This suggests that the selection notions under examination are relevant for "large" counterfactual changes. Future research can help clarify how to integrate "large" into a predictive model of selection.

9.   For further details on the effect of communication in repeated games with an unknown time horizon, see Fonseca and Normann (2012); Cooper and Kühn (2014); Harrington et al. (2016); Wilson and Vespa (2020).

the effects of explicit collusion, and these limits are predictable by theory. Using a change to the payoff primitives (here the discount rate), we make collusion a knife-edge, nonrobust equilibrium, and show that the effects of communication dissipate entirely.

While the experimental literature on repeated games has largely focused on the standard two-player RPD, there is a large literature studying a canonical $N$-player social dilemma: the voluntary contribution public-goods game (see Vesterlund, 2016, for a survey). Although much of this literature focuses on finite implementations, one notable exception is Lugovskyy, Puzzello, Sorensen, Walker, and Williams (2017). Similar to our paper, the authors use experimental variation over both $N$ and the payoff primitives (in their case, the return to the group contribution). However, this is done with a different end goal: to identify the isolated effect of the stage game's marginal per capita return. Instead, our objective is to isolate strategic uncertainty and test a predictive theory of selection.[10]

Beyond social dilemmas, our paper is also related to the literature on coordination games (see Devetag and Ortmann, 2007, for a survey). The strategic-uncertainty measure examined in our paper works because the RPD has a stag-hunt normal-form representation (Blonski and Spagnolo, 2015), adapting the risk-dominance notion for one-shot coordination games as in Harsanyi and Selten (1988).[11] Risk dominance has been shown to have substantial predictive content in simple coordination games with tradeoffs over payoff dominance and risk dominance (see Battalio, Samuelson, and Van Huyck, 2001; Brandts and Cooper, 2006; Dal Bó, Fréchette, and Kim, 2021, and references therein). Therefore, strategic uncertainty has demonstrated its usefulness as a theoretical selection device in both static and repeated games. We contribute to this literature with an experiment that explicitly tests and shows how the predictive effects extend further to multiplayer infinite-horizon settings.

Finally, our last robustness treatment provides a connection to coordination games with asymmetric-payoffs equilibria (such as the battle-of-the-sexes game). Coordination in this treatment requires at least two out of four players to cooperate, which relaxes the condition for success relative to the two-from-two treatment. But efficient equilibrium outcomes have two players coordinate on defecting (and getting a higher payoff) and two players cooperating (and getting a relatively lower payoff). The asymmetry means that each player would prefer to be a free-rider. Similar tensions

---

10.   Relatedly, Martinez-Martinez and Normann (2024) study an $N$-player social dilemma in continuous time and finds that as $N$ increases (and strategic uncertainty increases) cooperation decreases.

11.   The difference in our setting is that neither total payoffs nor strategic choices are directly provided to the participants, as these are extensive-form objects. Instead, participants are given the stage-game payoffs and actions, from which strategies (e.g., grim trigger or tit for tat) and gross payoffs are endogenously derived. Our use of risk dominance in a repeated game refers to the concept constructed by Spagnolo and Blonski (2001) inspired by Harsanyi and Selten (1988).

arise in one-shot coordination games with asymmetric payoffs like the battle-of-the-sexes game, where the literature has documented relatively high coordination failure rates that result in lower payoffs (cf. Cooper, DeJong, Forsythe, and Ross, 1990, 1993, 1994; Straub, 1995; Crawford, Gneezy, and Rottenstreich, 2008). The low cooperation rates in our two-from-four robustness treatment suggest that coordination challenges introduced by asymmetric payoffs already documented in the one-shot battle-of-the-sexes game can extend to a repeated-game setting like ours.[12]   try

## 2. Generalizing the Basin of Attraction

We begin this section by summarizing the progress made towards validating the basin of attraction for always defect as a theoretical prediction in the two-player RPD literature. A reader familiar with the literature can skip to Section 2.2, where we extend the framework by introducing a new parameter for strategic uncertainty, the number of players $N$.

### 2.1. Two-players

Consider an RPD with a discount rate $\delta \in (0,1)$. In each period $t = 1, 2, \ldots$ players $i \in \{1, 2\}$ simultaneously select actions $a_i \in \mathcal{A} := \{$(C)ooperate, (D)efect$\}$. The period-payoff for player $i$ is a function of both players' choices, $\pi_i(a_i, a_j)$, where all symmetric PD stage-games can be expressed in a compact form by normalizing all payoffs relative to the joint-defection payoff $\pi_0 := \pi(D,D)$, and rescaling with the relative gain from joint cooperation: $\Delta\pi := \pi(C,C) - \pi_0$.[13] Defining scale and normalization in this way, the PD stage-game can be expressed with two parameters $g$ and $s$ for the different-action payoffs $\pi_i(D,C) = \pi_0 + (1+g)\Delta\pi$ and $\pi_i(C,D) = \pi_0 - s\Delta\pi$. The parameters $g > 0$ and $s > 0$ capture the relative temptation- and sucker-payoffs, respectively.

The strategic-uncertainty measure we focus on is based on two focal extensive-form RPD strategies:[14]

1. *always defect*, $\alpha_{\text{All-D}}$, which plays the stage-game Nash in all rounds (the worst-case subgame-perfect equilibrium of the game).

---

2. *Grim trigger*, $\alpha_{\mathrm{Grim}}$, which begins by cooperating, but switches to always defect after any defection in past play (the collusive subgame-perfect equilibrium).[15]

As functions of the observable history $h_t$, these two strategies are given by:

$$\alpha_{\mathrm{Grim}}(h_t) = \begin{cases} C & \text{if } t = 1 \text{ or } h_t = ((C,C),(C,C),\dots,(C,C)), \\ D & \text{otherwise}; \end{cases}$$

$$\alpha_{\mathrm{All\text{-}D}}(h_t) = D.$$

Strategic uncertainty in the two-player RPD is measured through the size of the basin of attraction for always defect. The model considers the expected reward for player $i$ when uncertainty on the other player $j$ is represented by a believed strategy mixture $p \cdot \alpha_{\mathrm{Grim}} \oplus (1-p) \cdot \alpha_{\mathrm{All\text{-}D}}$. The *basin for always defect* is defined as the set of beliefs $p$ for which player $i$ receives a higher expected payment from $\alpha_{\mathrm{All\text{-}D}}$ than $\alpha_{\mathrm{Grim}}$. The always-defect belief basin is therefore the interval $[0, p^\star(g,s,\delta)]$ with the critical-point/interval-width given by:[16]

$$p^\star(g,s,\delta) \equiv \frac{(1-\delta)s}{\delta - (1-\delta)(g-s)} \tag{1}$$

Consequently, the PD stage-game payoffs are used as primitive inputs into a risk/reward model of collusion based upon strategic uncertainty.

Equation (1) represents a theoretical relationship between the payoff primitives of the game and a critical strategic belief over the other player's likelihood of collusion. The hypothesized relationship is monotone, which allows unambiguous directional predictions for any counterfactual change in the primitives. Moreover, the cardinal basin-size measure directly implies the ordinal risk-dominance relationship between the two strategies. If $p^\star(g,s,\delta) < 1/2$ the collusive strategy $\alpha_{\mathrm{Grim}}$ risk dominates $\alpha_{\mathrm{All\text{-}D}}$, and vice versa.

Using results from the meta-study on the two-player RPD (Dal Bó and Fréchette, 2018), we illustrate the relationships between the scalar basin-size measure of strategic uncertainty and our two focal outcome measures: initial and ongoing cooperation rates. In both panels of Figure 1, the horizontal axis represents the theoretical measure of strategic uncertainty, while the vertical axes represent one of our outcome measures. In Panel (A) we present the results for *initial cooperation*; in Panel (B) we present results for *ongoing cooperation*. The solid line in both panels indicates $\hat{\mathcal{C}}_{\mathrm{Meta}}(p^\star)$,

---

15. The strategy here is 'best case' as: (i) It can support the best-case outcome. (ii) It uses the harshest possible punishment and can support collusion at smaller values of $\delta$ than any other strategy. (iii) Any realized miscoordination is minimal and resolves in a single round.

16. In the case that the strategy $(\alpha_{\mathrm{Grim}}, \alpha_{\mathrm{Grim}})$ is not a subgame-perfect equilibrium of the repeated game, the basin size for always defect is defined as $p^\star(g,s,\delta) = 1$.
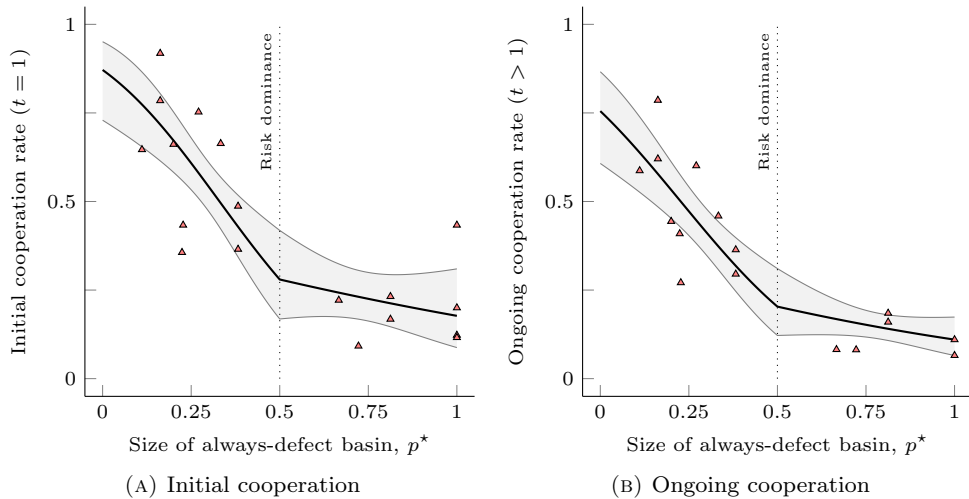
FIGURE 1. Meta-study relationship: strategic uncertainty and RPD cooperation—black line shows estimated relationship from (Dal Bó and Fréchette, 2018) meta-study (95 percent confidence in gray), with each point representing a separate treatment.

which we use to denote the predicted cooperation rate using meta-study data at each $p^\star$.[17] The shaded region represents the 95 percent confidence interval for $\hat{\mathcal{C}}_{\mathrm{Meta}}(p^\star)$.

For both initial and ongoing cooperation, we find essentially the same predicted relationship $\hat{\mathcal{C}}_{\mathrm{Meta}}(p^\star)$, consistently low levels of cooperation when always-defect is risk dominant ($p^\star > 1/2$); and a significantly decreasing relationship with $p^\star$ when collusion is risk dominant ($p^\star < 1/2$).

The theoretical model used in the basin construction posits a connection between initial and ongoing cooperation. If collusion functions through conditional cooperation with grim-trigger punishments, the expected ongoing cooperation rate is the probability that the players jointly cooperate in the first round: the initial cooperation rate squared. Thus, if cooperation were effectively governed by grim triggers, both measures of empirical cooperation would carry the same information. Since, in fact, grim-trigger punishments have been documented to be used by subjects (for example, Dal Bó and Fréchette, 2011), data from RPD games do not provide enough variation to identify whether theoretical notions track more closely with either empirical

---

17.    We estimate a probit regression using meta-study data from 996 participants clustered across 18 experimental treatments, where we focus on late-session cooperation (supergames 16-20). The individual-level cooperation decisions serve as the left-hand side variable, and the basin size is included on the right-hand side in a piecewise-linear fashion around the risk-dominance dividing point. Our econometric specification is inspired by Dal Bó and Fréchette (2018, Table 4). However, to maintain a continuous relationship, we modify their specification by eliminating a degree of freedom that allowed for a discontinuity at $p^\star = 1/2$.

measure. Consequently, with only two players it is challenging to identify the extent to which the strategic-uncertainty measure predicts initial intentions versus successful coordination.[18] However, as we will show below, adding more players provides additional variation that will allow us to differentiate between the two measures of cooperation.[19]

### 2.2. Extending to N > 2

We now extend the strategic-uncertainty model to an $N$-player environment. To achieve this, we consider a family of symmetric social dilemmas that nest the standard two-player RPD. To maintain a constant $2 \times 2$ stage-game representation for all $N$, our family of dilemmas makes use of an aggregate signal of the other agents' actions. All players $i = 1, \ldots, N$ continue to make a binary action choice $a_i \in \mathcal{A} \equiv \{C, D\}$, but their payoffs do not vary with (and they do not receive feedback on) the separate actions of the other $N - 1$ players. Instead, players' payoffs are determined by their own-action $a_i$ and a deterministic binary signal $\sigma(a_{-i}) \in \{S(\text{uccess}), F(\text{ailure})\}$ of the others' actions, $a_{-i}$. In particular, the generic player $i$'s stage-game payoff and signal function are given, respectively, by:

$$\pi_i\left(a_i, \sigma\right) = \begin{cases} \pi_0 + \Delta\pi & \text{if } a_i = C, \sigma = S, \\ \pi_0 + \Delta\pi\left(1 + x\right) & \text{if } a_i = D, \sigma = S, \\ \pi_0 - \Delta\pi x & \text{if } a_i = C, \sigma = F, \\ \pi_0 & \text{if } a_i = D, \sigma = F; \end{cases} \quad (2)$$

$$\sigma(a_{-i}) = \begin{cases} S & \text{if } a_j = C \text{ for all } j \neq i, \\ F & \text{otherwise.} \end{cases} \quad (3)$$

These choices lead to a symmetric game, in which payoffs can be summarized with a $2 \times 2$ table over: (i) The own action $C$ or $D$; and (ii) The signal outcome, an $S$ signal if the other $N - 1$ players jointly cooperate, or an $F$ signal if at least one other player defects.[20]

Ignoring the scale and normalization of the game (held constant in our experiments with $\Delta\pi = \$9$ and $\pi_0 = \$11$), the repeated games we examine are summarized by

---

18.  For a setting that achieves this with sequentiality of moves, see Ghidoni and Suetens (2022).

19.  In settings where collusion requires $N$ agents to initially cooperate to produce ongoing cooperation, the relationship is given by initial cooperation rate to the $N$-th power. Separate identification between the two measures is possible by comparing treatments with different values of $N$.

20.  In Section 5, we present a treatment where cooperative outcomes require only two out of four players to cooperate. Introducing the possibility of achieving cooperative outcomes with some players not cooperating gives rise to a free-riding problem. However, in our main treatments, we sidestep this issue by assuming efficient ongoing cooperative outcomes only when all $N$ players cooperate.

three primitives: (i) The relative cost of cooperating, $x$;[21] (ii) The number of players, $N$; and (iii) The continuation probability, $\delta$. Our experiments fix $\delta = 3/4$ in all but one diagnostic treatment in Section 5. This leaves us with two key experimental parameters: the relative cost of cooperating $x$ and the number of participants $N$.

In building a model of strategic uncertainty for arbitrary $N$, we use a symmetric belief over the others' choices. That is, we assume each player chooses a mixture over $\alpha_{\text{Grim}}$ and $\alpha_{\text{All-D}}$.[22] Our family of social dilemmas requires cooperation from all $N$ players for everyone to get an $S$ signal. Thus, strategic uncertainty reduces to the probability that the other $N - 1$ players *jointly* coordinate on the collusive strategy,

$$Q(N) = \Pr\{N - 1 \text{ others all choose } \alpha_{\text{Grim}}\}.$$

In every other case, at least $N - 1$ players will receive an $F$ signal and the punishment path will be triggered.

As in the case of two players, the critical belief $Q^\star(N)$ is given by the point of indifference between the amount given up with certainty from a single round of cooperation, $x\Delta\pi$, and the continuation gain from collusion, $[\delta/(1 - \delta)] \cdot \Delta\pi$, obtained with probability:

$$Q^\star(N) = \frac{(1 - \delta)}{\delta}x,$$

where the right-hand-side is identical to the two-player construction in Equation (1) for $x = g = s$.

Next, we need to relate the joint cooperation of the other $N - 1$ players to the probability $p$ that each individual other player attempts to collude. Our design focuses on two extremes. The "standard" extension in which beliefs are fully independent; and an alternative/null-effect model in which beliefs are perfectly correlated.[23]

Assuming perfect correlation for the other $N - 1$ agents, $Q(N) = p$, so the critical belief is:

$$p_{\text{Corr.}}^\star(x) = \frac{1 - \delta}{\delta}x. \tag{4}$$

---

21. In the meta-study notation this is implemented with $s = g = x$. This single-parameter formulation is equivalent to the Fudenberg, Rand, and Dreber (2012) benefit/cost formulation, where their benefit/cost ratio parameter $(b/c)$ is given by $(1 + x)/x$ here.

22. For the $N$-player dilemma we define the grim-trigger strategy with imperfect signals as:

$$\alpha_{\text{Grim}}(h_t) = \begin{cases} C & \text{if } t = 1 \text{ or } h_t = ((C, S), (C, S), \ldots, (C, S)), \\ D & \text{otherwise.} \end{cases}$$

23. See Cason, Sharma, and Vadovič (2020) for an example of correlated beliefs that emerge in situations where independence would be the standard prediction.

In contrast, when beliefs are fully independent, $Q(N) = p^{N-1}$, so the critical belief is:

$$p_{\text{Ind.}}^{\star}(x, N) = \left(\frac{1-\delta}{\delta} x\right)^{1/N-1} \equiv \left(p_{\text{Corr.}}^{\star}(x)\right)^{1/N-1}. \tag{5}$$

Note that the correlated measure in Equation (4) is not a function of $N$, while the independent measure in Equation (5) increases in $N$. The two measures are identical only in the RPD case at $N = 2$.[24]

We focus on these two extreme cases of full independence and perfect correlation because: (i) They allow us to produce an experimental design that has stark behavioral predictions; and (ii) Both measures are simple to compute in settings beyond our environment.[25]

## 3. Experimental Design

Based on the basin measures derived in Equations (4) and (5), our experimental design is founded on two competing hypotheses:

CORRELATED-BASIN/NULL-EFFECT HYPOTHESIS. *Cooperation decreases as we increase the cost of cooperation $x$, but there is no effect as we vary the number of players $N$.*

INDEPENDENT-BASIN HYPOTHESIS. *Cooperation decreases as we increase $x$ and/or $N$. Moreover, the substitution effects between $x$ and $N$ indicate no effect on cooperation if we decrease $x$ and increase $N$ to hold constant $p_{Ind.}^{\star}$.*

In Panel (A) of Table 1 we illustrate our first treatment dimension, which manipulates the payoff cost of cooperating $X = x\Delta\pi$, where $\Delta\pi = \$9$. The two values of $X$—a high temptation of \$9 (a normalized temptation of $x = 1$) illustrated on the left,

---

24.    Notice that both extensions of the measure capture beliefs over supergame strategies (full specifications of what action to play at any history). For the two strategies underlying the basin measures, actions are perfectly correlated in all rounds after the first one. For instance, consider $\alpha_{\text{Grim}}$. Either all $N$ players successfully coordinate on cooperation, or after an observed failure in round one, the punishment path is triggered with all $N$ players choosing defect in all subsequent rounds. As such, the independent and correlated models will only differ in the potential for correlation in the very first round.

25.    One can define an intermediate hypothesis with an extra parameter that captures the extent to which beliefs are independent (with complementary probability on the extent to which beliefs are correlated). In Section 4 we discuss this alternative in further detail and estimate the correlation parameter from the data.

TABLE 1. Experimental design

| **Panel A.** | **Stage-game payoffs** | $X = \$9$ | | $X = \$1$ | |
|---|---|---|---|---|---|
| | | $\sigma(a_{-i}) = S$ | $\sigma(a_{-i}) = F$ | $\sigma(a_{-i}) = S$ | $\sigma(a_{-i}) = F$ |
| Cooperate, $\pi_i(C,\sigma)$ | | $\$20$ | $\$2$ | $\$20$ | $\$10$ |
| Defect, $\pi_i(D,\sigma)$ | | $\$29$ | $\$11$ | $\$21$ | $\$11$ |

| **Panel B.** | **All-D Basin Size** | $X = \$9$ $(x = 1)$ | | $X = \$1$ $(x = 1/9)$ | |
|---|---|---|---|---|---|
| | | $N = 2$ | $N = 4$ | $N = 4$ | $N = 10$ |
| Correlated, $p^\star_{\text{Cor.}}(x)$ | | $p^\star_0$ [0.33] | $p^\star_0$ [0.33] | $p^\star_0 - \Delta p^\star_{\text{Cor.}}$ [0.04] | $p^\star_0 - \Delta p^\star_{\text{Cor.}}$ [0.04] |
| Independent, $p^\star_{\text{Ind.}}(x,N)$ | | $p^\star_0$ [0.33] | $p^\star_0 + \Delta p^\star_{\text{Ind.}}$ [0.69] | $p^\star_0$ [0.33] | $p^\star_0 + \Delta p^\star_{\text{Ind.}}$ [0.69] |
| Sessions | | 3 | 3 | 3 | 2 |
| Subjects | | 60 | 60 | 72 | 60 |

| **Panel C.** | **Meta-study prediction** | $p^\star_0$ | Marginal effect from basin: | |
|---|---|---|---|---|
| | | [0.33] | Increase to [0.69] | Decrease to [0.04] |
| Initial coop., $t = 1$ | | 0.50 | $-0.26$ | $+0.35$ |
| Ongoing coop., $t > 1$ | | 0.37 | $-0.21$ | $+0.35$ |

Meta-study predictions in Panel (C) correspond to the estimated relationship $\hat{\mathcal{C}}_{\text{Meta}}(p^\star)$ illustrated in Figure 1 .

and a low temptation of \$1 (a normalized temptation of $x = 1/9$) illustrated on the right—lead to two payoff environments over own actions and signals.[26,27]

We also vary the number of players $N$ as indicated in the column headings of Panel (B) in Table 1. The two rows of Panel (B) illustrate how choices regarding $X$ and $N$ influence the basin-size measures of strategic uncertainty under the correlated and independent extensions. In total, we create four treatments, each defined by an $(N, X)$-pair.

To independently manipulate each basin-size measure, we select $(N{=}2; X{=}\$9)$ as our baseline treatment. When comparing $(N{=}2; X{=}\$9)$ with $(N{=}4; X{=}\$9)$ we keep the correlated-basin measure constant at $p^\star_0 = 0.33$ and increase the independent-basin measure to $p^\star_0 + \Delta p^\star_{\text{Ind.}} = 0.69$. Next, when comparing $(N{=}2; X{=}\$9)$ with $(N{=}4; X{=}\$1)$ we keep the independent-basin measure constant at $p^\star_0 = 0.33$ and lower the correlated-basin measure to $p^\star_0 - \Delta p^\star_{\text{Corr.}} = 0.04$. Finally, when comparing $(N{=}4; X{=}\$1)$ with $(N{=}10; X{=}\$1)$ we keep the correlated-basin measure constant at $p^\star_0 - \Delta p^\star_{\text{Corr.}} = 0.04$ and increase the independent-basin measure to $p^\star_0 + \Delta p^\star_{\text{Ind.}} = 0.69$.

---

26.  See Figure E.1 in Online Appendix E for representative lab screenshots.

27.  Henceforth, we will focus on the payoff cost of cooperating $X$ rather than the normalized parameter $x$.

By varying the primitives $X$ and $N$, our $2 \times 2$ design yields four pairs of correlated/independent basin measures:[28]

$$(p_{\text{Corr.}}^\star, p_{\text{Ind.}}^\star) \in \quad \left\{ p_0^\star, p_0^\star - \Delta p_{\text{Corr.}}^\star \right\} \times \quad \left\{ p_0^\star, p_0^\star + \Delta p_{\text{Ind.}}^\star \right\}$$
$$:= \left\{ 0.33, 0.04 \right\} \times \quad \left\{ 0.33, 0.69 \right\}.$$

Using the probit-model estimates illustrated in Figure 1 we can provide a *quantitative* prediction $\hat{\mathcal{C}}_{\text{Meta}}(p^\star)$ for the cooperation rate under each basin-size measure $p^\star$. These predictions are outlined in Panel (C) of Table 1. The first column presents the initial and ongoing cooperation rates expected at $p^\star = 0.33$. The next two columns indicate the expected treatment effect resulting from a shift in strategic uncertainty from $p^\star = 0.33$ to either $p_0^\star - \Delta p_{\text{Corr.}}^\star = 0.04$ or $p_0^\star + \Delta p_{\text{Ind.}}^\star = 0.69$.

For illustration, consider the predictions under the standard independence-based extension. In the ($N{=}2$; $X{=}\$9$) and ($N{=}4$; $X{=}\$1$) treatments, the independent basin size is 0.33, and it increases to 0.69 in ($N{=}4$; $X{=}\$9$) and ($N{=}10$; $X{=}\$1$). If the strategic uncertainty relationship estimated from the two-player RPD meta-data were perfectly extrapolated to our setting, we should expect: (i) A reduction of 26 (21) percentage points in initial (ongoing) cooperation across the treatment pairs, caused by an increase in strategic uncertainty. (ii) A null effect on cooperation within each treatment pair, reflecting the designed perfect substitution across $X$ and $N$ in the independence-based measure.[29]

Note that our hypotheses do not specify whether initial cooperation, ongoing cooperation, or both are expected to align with the behavior of many players. Initial cooperation captures intentions to coordinate (with beliefs as a driver), while ongoing cooperation reflects successful coordination (with the interaction of the beliefs as a driver). In the case of the two-player RPD, Figure 1 shows that the basin size closely follows both cooperation measures, making it challenging to disentangle the effects. By introducing variation in $N$, we add a channel that might help us distinguish between initial and ongoing cooperation, and identify which measure is better predicted by basin-size models.

---

28.   We chose $\Delta \pi = \$9$ and $\delta = 3/4$ for simplicity of presentation to participants (i.e., integer values for both $N$ and $X$). The precise design over the basin measures is as follows:

$$(p_{\text{Corr.}}^\star, p_{\text{Ind.}}^\star) \in \left\{ 3^{-1}, 3^{-3} \right\} \times \left\{ 3^{-1}, 3^{-1/3} \right\}.$$

29.   Alternatively, under a null-effect from $N$, given by the correlated-basin measure, the basin size is reduced from 0.33 to 0.04 as we move between the ($N{=}2$; $X{=}\$9$) and ($N{=}4$; $X{=}\$9$) treatment pair and the ($N{=}4$; $X{=}\$1$) and ($N{=}10$; $X{=}\$1$) pair. Based on the estimated relationship from the meta-study, this implies an increase in the initial cooperation rate of 35 percentage points and an increase in the ongoing cooperation rate of 50 percentage points, and null effects within each pair for fixed $X$.

### Experimental Specifics

In our main experiments, we used a between-subject design over the four treatments outlined in Table 1. Participants for each treatment were recruited from the undergraduate population at the University of Pittsburgh, and each took part in exactly one session. We recruited a total of 584 participants, 252 for the first four main treatments, and 332 for the extensions discussed in Section 5. Each treatment comprised three sessions, aiming to enroll a minimum of 20 participants per session, except for the $(N{=}10; X{=}\$1)$ treatment, for which we conducted two sessions with 30 participants each.[30] Sessions lasted between 55 and 90 minutes, and participants received an average payment of approximately \$19.

Each session comprised 20 supergames, with a common random termination chance of $1 - \delta = 1/4$ after each completed round.[31] The participants were randomly and anonymously matched in the 20 supergames in a stranger design.[32] The 20 supergames were divided into two parts of ten supergames.[33] For final payment, one supergame from each part was randomly selected, where only the actions/signals from the last round in the selected supergame counted for payment.[34]

## 4. Results

We begin this section by describing the aggregate cooperation and success rates at the treatment level. Then, we proceed to discussing inferential tests of our two basin-extension hypotheses.

---

30.    While our design called for sessions to have at least 20 participants, we allowed sessions to grow by an additional group of size $N$ depending on realized show ups. For $(N{=}10; X{=}\$1)$ we instead opted to recruit 30 participants for each session so that we had three groups in each supergame.

31.    We employed common draws to maintain consistent supergame lengths at the session level for each treatment.

32.    All participants received both written and verbal instructions regarding the task and payoffs. Detailed instructions are available for interested readers in the Online Appendix F.

33.    Participants were provided with complete instructions for the first part and were informed that instructions for the second part would be given after completing supergame ten. For the four between-subject treatments outlined in Section 3, part two was identical to part one. In later sections of the paper, we describe an additional set of treatments with a within-subject change across the two parts. The decision to have two identical parts here enables direct comparisons in first-half play.

34.    This method, developed in Sherstyuk, Tarui, and Saijo (2013), is employed to induce risk neutrality across supergame lengths. Another advantage of this design choice is the absence of wealth effects within a supergame, where history serves only as an instrument for the future play of others.

TABLE 2.  Cooperation and success rates across all supergames

| Action and signal rates | $X = \$9$ | | $X = \$1$ | |
|---|---|---|---|---|
| | $N = 2$ | $N = 4$ | $N = 4$ | $N = 10$ |
| | Cooperation | | | |
| Initial | 0.503 (0.058) | 0.035 (0.017) | 0.792 (0.042) | 0.357 (0.055) |
| | $\langle 0.50 \rangle$ $\lvert 0.50 \rvert$ | $\langle 0.24 \rangle$ $\lvert 0.50 \rvert$ | $\langle 0.50 \rangle$ $\lvert 0.84 \rvert$ | $\langle 0.24 \rangle$ $\lvert 0.84 \rvert$ |
| Ongoing | 0.450 (0.055) | 0.006 (0.003) | 0.409 (0.050) | 0.185 (0.048) |
| | $\langle 0.37 \rangle$ $\lvert 0.37 \rvert$ | $\langle 0.16 \rangle$ $\lvert 0.37 \rvert$ | $\langle 0.37 \rangle$ $\lvert 0.72 \rvert$ | $\langle 0.16 \rangle$ $\lvert 0.72 \rvert$ |
| | Success | | | |
| Initial | 0.503 | 0.000 | 0.578 | 0.000 |
| Ongoing | 0.450 | 0.000 | 0.293 | 0.000 |

Results are calculated using data from the last-five supergames. Cooperation rates present raw proportions, with subject-clustered standard errors in parentheses. For comparison, we provide the meta-study prediction for the independent basin measure $\hat{\mathcal{C}}_{\text{Meta}} \left( p^{\star}_{\text{Ind.}} \right)$ in angle-brackets, $\langle \cdot \rangle$, and the prediction $\hat{\mathcal{C}}_{\text{Meta}} \left( p^{\star}_{\text{Cor.}} \right)$ for the correlated basin measure in vertical bars, $\lvert \cdot \rvert$ (cf. Panel (B) of Table 1 for details).

## 4.1. Main Treatment Differences

In Table 2 we present average cooperation and success rates by treatment, for both initial and ongoing cooperation. Averages are computed for the last five supergames to capture late-session behavior, where subjects have accumulated experience in the environment.[35] Overall, the results reveal large shifts in cooperation as we manipulate the cost of cooperation $X$ and/or the number of players $N$.

The first row in Table 2 provides a summary of initial cooperation. The 50.3 percent initial cooperation rate in our ($N=2$; $X=\$9$) treatment closely aligns with the 50.0 percent rate predicted by the meta-study. However, maintaining the cooperation cost at $X = \$9$ and doubling the group size to four virtually eliminates cooperative behavior, resulting in an initial cooperation rate of 3.5 percent in ($N=4$; $X=\$9$). In the first round of our low-temptation scenarios ($X = \$1$), groups of $N = 4$ exhibit highly cooperative behavior (79.2 percent) while groups of $N = 10$ display moderate cooperation (35.7 percent).

The next two rows in Table 2 summarize the ongoing cooperation rates. Across all treatments we observe a decrease in ongoing cooperation compared to initial cooperation. The most substantial quantitative drops are evident in the $X = 1$

---

35. Including all rounds yields similar results (see Table A.1 in Online Appendix A).

treatments, where ongoing cooperation rates are halved in comparison to the initial cooperation rates.[36]

The last two rows in Table 2 present the fraction of rounds in which a success signal was observed.[37] The patterns that emerge for success rates are similar to those seen for ongoing cooperation, though with starker quantitative effects. Although success is the modal signal in the ($N$=2; $X$=\$9) and ($N$=4; $X$=\$1) treatments, in the ($N$=4; $X$=\$9) and ($N$=10; $X$=\$1) treatments there are *no* successes at all.[38]

The results presented in Table 2 speak to both the correlated- and independent-basin hypotheses. The collected evidence does not favor the correlated-basin hypothesis. For both the initial and ongoing cooperation rates we observe large changes in behavior as we move $N$ for either fixed value of $X$. On the other hand, the data support the independent-basin predictions regarding directional shifts in both initial and ongoing cooperation rates as we vary $X$ or $N$ in isolation. However, for initial cooperation, we observe deviations from perfect substitution of strategic uncertainty as $X$ and $N$ move in opposing directions. The independent-basin hypothesis predicts a null effect when we compare either ($N$=2; $X$=\$9) to ($N$=4; $X$=\$1) or ($N$=4; $X$=\$9) to ($N$=10; $X$=\$1). But instead we observe substantial effects in the comparisons of initial cooperation, with 29 and 35 percentage point differences, respectively.

Meanwhile, ongoing cooperation in the ($N$=2; $X$=\$9) and ($N$=4; $X$=\$1) treatments are relatively close, at 45.0 and 40.0 percent, respectively. This finding aligns qualitatively with the independent-basin prediction of no difference due to perfect substitution of strategic uncertainty. For a similar comparison, however, we still note an 18 percentage point difference between ($N$=4; $X$=\$9) and ($N$=10; $X$=\$1). The difference is driven by a very stark finding of near-zero cooperation in ($N$=4; $X$=\$9). As we outline further below this is the main deviation in our data relative to the meta-study prediction.[39]

---

36.    In Online Appendix A, Table A.2 provides a more detailed breakdown of ongoing cooperation rates based on the observed history from the previous round. The findings suggest that individual cooperation is heavily conditioned on successful coordination in the preceding round. Interestingly, participants are markedly more forgiving after failed cooperation at $X = \$1$ than $X = \$9$.

37.    A success at the individual level requires joint cooperation from the other $N - 1$ participants. Success is directly linked to group-level cooperation, where the *expected* success rate, given an independent cooperation rate $p$, is $p^{N-1}$. In two-player games, the success rate is identical to the cooperation rate. Expected success rates (given the cooperation rate and independent matching) in the initial round are, in the order of Table 2 columns: 0.503, $4.2 \times 10^{-5}$, 0.497 and $9.3 \times 10^{-5}$.

38.    As success directly aggregates individual-level cooperation, we refrain from reporting standard errors (where standard errors also cannot be calculated in cases where we have no variation). Nevertheless, the pronounced nature of the effect, in alignment with predictions for the independent-basin hypothesis, clearly illustrates the underlying economic relationship.

39.    Regarding inference, Online Appendix A presents two additional tests: (i) Tests examining the cardinal predictions from the meta-study, and (ii) Tests assessing the ordinal predictions across treatments. In the first set of tests (Table A.3), we evaluate the predicted cooperation levels $\hat{\mathcal{C}}_{\text{Meta}}(p^{\star})$ from the meta study. Our findings reveal the rejection of cardinal predictions for both initial and ongoing cooperation, irrespective of whether the basins are independent or correlated ($p < 0.001$ all $F$-tests). However, closer examination indicates that the meta-study aligns more

TABLE 3. Basin-effect decomposition: Main treatments

| Experimental results | $p_0^\star$ | Marginal effect in cooperation from: | |
|---|---|---|---|
| | [0.33] | Independent basin increase to $p_0^\star + \Delta p_{\text{Ind.}}^\star = [0.69]$ | Correlated basin decrease to $p_0^\star - \Delta p_{\text{Corr.}}^\star = [0.04]$ |
| Initial | 0.464 (0.051) ⟨0.50⟩ | −0.395 (0.048) ⟨−0.26⟩ | +0.357 (0.053) ⟨+0.35⟩ |
| Ongoing | 0.366 (0.051) ⟨0.37⟩ | −0.293 (0.051) ⟨−0.21⟩ | +0.115 (0.061) ⟨+0.35⟩ |

Results are calculated using data from the last-five supergames. The cooperation decomposition runs two probits, one for initial, and one for the ongoing cooperation, with subject-clustered standard errors in parentheses. Right-hand-side variables are a constant and two dummies, one for a low-correlated–basin treatment ($X = \$1$, both $N$ values), one for a high-independent–basin treatment ($X = \$9/N = 4$ and $X = \$1/N = 10$). Meta-study predictions given in angle brackets, $\langle \cdot \rangle$, below each result.

## 4.2. Evaluation of the Independent- and Correlated-Basin Hypotheses

Table 3 presents a direct statistical evaluation of our two competing hypotheses. The results are based on probit regressions that examine subjects' cooperation decisions, with dummy variables corresponding to the 2×2 design outlined in Table 1 Panel (B). The dummy covariates include an indicator for the predicted $\Delta p_{\text{Cor.}}^\star$ decrease from the correlated basin (as we decrease $X$ for any $N$) and an indicator for the predicted $\Delta p_{\text{Ind.}}^\star$ increase from the independent basin (as we increase $N$ holding $X$ constant).

Each row in Table 3 presents results from a distinct estimation, one focusing on initial cooperation and the other on ongoing cooperation. The first $p_0^\star$ column displays the estimated cooperation rate when both dummy variables are zero, representing the RPD cooperation rate with a basin size of $p_0^\star = 0.33$. The following two columns illustrate the estimated marginal effect on the cooperation rate for a shift in each basin measure, while holding the other basin constant. If either of the two basin hypotheses fully explained behavior, we would expect a significant estimate for the dummy on that basin shift and an insignificant effect on the other.

---

closely with predicted behavior in two specific scenarios: (i) ongoing cooperation, and (ii) predictions using the independent basin. For ongoing cooperation under the independent-basin prediction, we do not reject the predicted cooperation levels in ($N$=2; $X$=\$9), ($N$=2; $X$=\$1) and ($N$=10; $X$=\$1) ($p > 0.150$ all comparisons, jointly $p = 0.447$). The only exception is the ($N$=4; $X$=\$9) treatment ($p < 0.001$), where the meta-study predicts a cooperation rate of 16 percent, but the observed rate is virtually zero.

For the ordinal comparisons, Table A.4 presents the six possible treatment comparisons in our design, along with the ordinal prediction from each basin notion. For ongoing (initial) cooperation, the independent basin correctly organizes five (four) of the six comparisons. While it orders the four treatments where a difference is predicted (for both ongoing and initial cooperation), the independent basin fails one of the two null tests for ongoing cooperation and both null tests for initial cooperation. Meanwhile, the correlated basin makes three successful predictions out of six, one incorrect directional prediction, and fails both null tests.
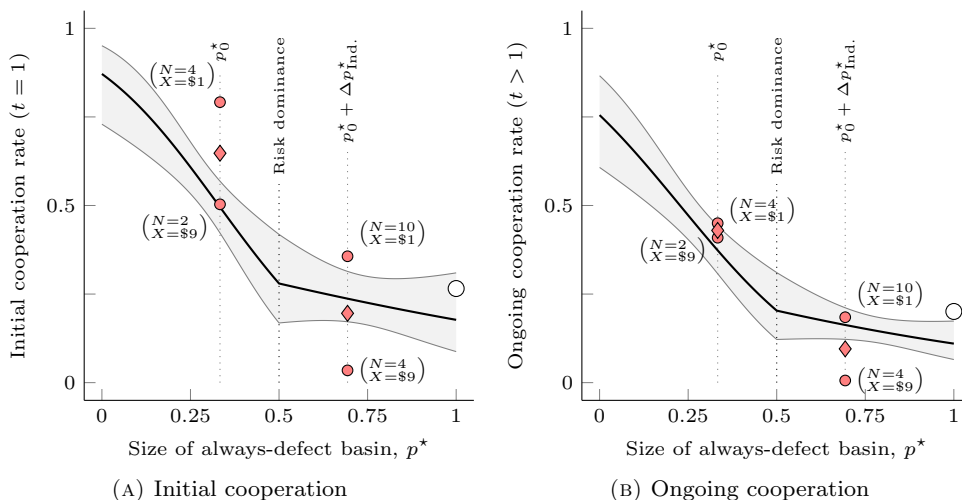
FIGURE 2. Cooperation under the independent basin-size model—filled circles indicate treatment averages; filled diamonds pooled average over each independent-basin value; empty circle indicates unilateral cooperation rates from extension treatment discussed in Section 5.

The estimation parallels the probit model we run to recover the meta-study prediction $\hat{\mathcal{C}}_{\text{Meta}}(p^{\star})$. The estimated baseline cooperation rates is for an RPD with $p_0^{\star} = 0.33$, where this baseline closely matches the meta-study prediction. Specifically, while the meta-study predicts the initial (ongoing) cooperation rate of 50.0 (37.0) percent, our data at $p_0^{\star} = 0.33$ reflects a very similar (and statistically indistinguishable) rate of 46.4 (36.6) percent. To illustrate this alignment, Figure 2 depicts the fitted relationships from the meta-study overlaid with our results from the four treatments using the independent-basin size on the horizontal axis. Filled circles represent individual treatments and filled diamonds treatments pooled over each value for the independent-basin measure. While there is notable divergence for initial cooperation, Figure 2 demonstrates quantitatively similar results for ongoing cooperation.[40]

To test our two competing hypotheses, we focus on the second and third columns of Table 3. In the scenario where the independent-basin measure comprehensively captured all pertinent aspects of behavior, we would expect a statistically significant and negative estimate in the second column, coupled with a zero effect in the third column. Meanwhile, if the correlated-basin captured behavior, we would expect a significantly negative effect in the third column and a zero effect in the second column.

---

40. In Online Appendix A, Figure A.1 presents analogous results organized under the correlated-basin model. The figure illustrates much poorer organization of the data, both in terms of relative treatment comparisons and quantitatively.

In terms of initial cooperation, our results reveal that modifications to both basin measures yield significant effects ($p < 0.001$). Although the estimated effects are comparable in magnitude, they exhibit opposite directions, which is consistent with our predictions. Given that neither effect prevails over the other, we infer that both $X$ and $N$ contribute unique information to the prediction of initial cooperation, and this information is not entirely captured by either basin measure independently.

Regarding ongoing cooperation, the increase from the independent basin shift is negative and significant ($p < 0.001$) and it is quantitatively close to the meta-study prediction. Meanwhile the estimated effect for the correlated basin is much smaller in magnitude and significant only at the 10 percent level ($p = 0.061$) after controlling for the independent basin. The small effect attributed to the correlated basin in our estimation could also be associated with other-regarding preferences. Part of the differences in cooperation are driven by a higher fraction of unconditional cooperators at $X = \$1$ compared to $X = \$9$.[41] Because variation in $X$ is associated with shifts in the correlated-basin value (invariant to $N$), this presents a difficulty in interpretation for the small positive effects for the correlated basin. While this could be driven by belief correlation, it could also be driven by other-regarding preferences.[42]

In addition to the qualitative directional effects, we observe that the quantitative shifts in ongoing cooperation under the independent-basin measure closely align with the predicted effects expected from the meta-study.[43] Specifically, the latter predicts a drop of 21 percentage points (last row in Table 3) in ongoing cooperation when the size of the basin increases from $p_0^\star = 0.33$ to $0.69$, and our estimates indicate a decrease of 29 percentage points.[44]

---

41.    In Tables G.1 and G.2 of Online Appendix G we present strategy frequency estimates from the first and last seven supergames, where we identify the fraction of choices that are consistent with unconditional cooperation in each treatment.

42.    By design, subjects receive coarse feedback in our environment. For example, a failure signal in a treatment with $N = 4$ indicates that at least one of the other three members did not cooperate. Such coarse feedback minimizes the possibility that early feedback is exacerbated as $N$ increases. If we had provided subjects with disaggregated feedback, a treatment with $N = 10$ would provide effective feedback on everyone else in the session after a few supergames. This could translate into early choices having more of an impact in later choices in treatments with high $N$. While this type of effect is muted given our coarse feedback, it is still possible for it to arise. We do not see any clear evidence in this direction, but a definitive test would require treatments with a turnpike, perfect stranger designs, and/or larger sessions.

43.    Our measures of equilibrium selection aim to capture strategic uncertainty in a setting that differs from the two-player RPD. Discovering that our results align with findings in the two-player RPD literature is valuable. It implies that a measure of strategic uncertainty might be a robust predictor of collusion, irrespective of the specific details of the environment.

44.    On the contrary, the meta-study predicts an *increase* of 50 percentage points in initial cooperation when the size of the correlated basin decreases from $p_0^\star = 0.33$ to $0.04$, and our estimates indicate an increase of roughly 11.5 percentage points, after controlling for the independent-basin effects.

The main difference in ongoing cooperation between our data and the independent-basin predictions from the meta-study arises from extreme behavior in the ($N=4$; $X=\$9$) treatment, where cooperation is essentially at the boundary. A two-player RPD with a basin size of $p^\star = 0.69$ has a predicted ongoing cooperation rate of 16.0 percent, and this prediction remains relatively stable for all other values of the basin where grim is risk-dominated (with 11.0 percent cooperation predicted at $p^\star = 1$). The very low late-session cooperation rates in ($N=4$; $X=\$9$) can be explained by considering the large payoff reduction from cooperation, coupled with unrelentingly negative feedback. That is, out of 1,145 supergame-rounds in this treatment where a group of four attempted to coordinate, only a single group was successful for a single round. However, though the observed level deviates from the prediction, a broader interpretation of the basin continues to hold: conditional cooperation is not expected when always-defect is risk dominant ($p^\star > 1/2$), while the level of cooperation is predictably decreasing when grim is risk dominant ($p < 1/2$).

Finally, we attempt to measure *how much* correlation is necessary to rationalize the data. To achieve this, we allow beliefs to be a convex combination of the independent and correlated models. With proportion $\sigma$, the $N-1$ agents collectively choose grim with probability $p$ and always defect with probability $1-p$; with proportion $1-\sigma$, each agent independently chooses grim with probability $p$ and always defect with probability $1-p$. Under this specification, the probability that the other $N-1$ players coordinate is given by

$$\sigma \cdot p + (1-\sigma) \cdot p^{N-1},$$

with the critical belief denoted by $p^\star(\sigma, x, N)$. The additional parameter $\sigma$ nests the two extremes: $\sigma = 0$ for full independence, $\sigma = 1$ for perfect correlation.[45]

Looking at the best fitting parameter, for (I)nitial cooperation, we estimate $\hat{\sigma}_I = 0.091$ (with standard error of 0.005), while the comparable estimate for (O)ngoing cooperation is $\hat{\sigma}_O = 0.031$ (with the standard error of 0.014). Both estimates are statistically different from zero ($p < 0.001$ and $p = 0.014$ for initial and ongoing, respectively). The conclusion from the exercise is that the estimated degree of correlation needed is quantitatively small.

We now summarize our main results:

RESULT 1 (Independent-Basin Measure). *The independent-basin measure qualitatively organizes the results for ongoing cooperation and matches the quantitative level predictions in all treatments except for one. However, it does not contain all relevant information to predict initial intentions to cooperate.*

---

45.   Full details of the estimation procedure are provided in Online Appendix D.

RESULT 2 (Correlated-Basin Measure). *Our data are inconsistent with the predictions from the correlated-basin hypothesis, for both initial and ongoing cooperation. In particular, where the correlated basin predicts that behavior should, ceteris paribus, be unaffected by $N$, we find decreases in cooperation as $N$ increases. Quantitatively, the estimated degree of belief correlation required to rationalize the results is small.*

## 5. Beyond the main results

Our analysis thus far has abstracted away other features of the coordination problem to focus on the pure effects of the stage-game primitives. In this section, we introduce additional treatments to study possible limitations of the strategic-uncertainty model in predicting changes in equilibrium selection.

### 5.1. Between vs. Within Identification

Here, we explore the extent to which behavior after a policy change might not align with corresponding changes to the basin. Consider a policy change that alters the underlying strategic environment—the temptation and/or the number of players for our experiments—and therefore the collusive prediction. The underlying idea from the model is that beliefs about others' strategies drive behavior. But if beliefs are shaped by experiences formed prior to the policy change, a strategic uncertainty model might fail to predict behavioral changes *within* population. In the previous section, our treatments employed a *between*-subjects design, where identification relied on comparisons of late-session behavior between different populations, each with experience in a fixed strategic setting. In our modified treatments, we investigate the effects on collusive behavior following a change in the number of players $N$ *within* the same session.

We examine two within-session treatment shifts: one with ($N{=}2$; $X{=}\$9$) in the first half of the session, and ($N{=}4$; $X{=}\$9$) in the second half; and the reverse treatment shift with ($N{=}4$; $X{=}\$9$) in the first half, and ($N{=}2$; $X{=}\$9$) in the second. Given that we keep the temptation parameter constant at $X = \$9$, we label these two treatments as $2 \rightarrow 4$ and $4 \rightarrow 2$, respectively. In both treatments, the change in $N$ comes as a surprise: subjects are aware of a second part, but they do not receive instructions for the second part until the end of supergame ten. In terms of the independent-basin model, this creates a shift across the session from a low basin size of 0.33 when $N = 2$ to a high basin size of 0.69 when $N = 4$. In particular, this is a shift in $N$ that generates a substantial treatment effect for the between-subject treatments.

In Figure 3(A) we present the initial cooperation rates by supergame and type of treatment. The between-subject treatments with $N = 2$ and $N = 4$ are indicated by gray dashed lines, while the within-subject treatments are represented by two colored

lines: a solid red line for the $2 \rightarrow 4$ treatment and a dash-dotted blue line for the $4 \rightarrow 2$ treatment.

The figure illustrates a substantial between-subject effect, with more cooperation in $N = 2$ than $N = 4$ across all twenty supergames. Pooling the between and within treatments in supergames 6–10, we arrive at an initial cooperation rate of 47.4 percent for $N = 2$ and 13.9 percent for $N = 4$.[46] As we move into supergames 11–20 for our within-subject treatments, the strategic environment changes, specifically, the number of players an individual is matched with either decreases or increases. For the $2 \rightarrow 4$ treatment (the solid red line), initial cooperation remains high as $N$ increases. While there is no immediate drop in cooperation, we observe that as participants gain experience at $N = 4$, the cooperation rate continues to fall, reaching 16.7 percent by supergame 20. In contrast, moving from $N = 4$ to $N = 2$ (the blue dash-dotted line), we observe an immediate jump in cooperation as $N$ decreases: the initial-round cooperation in the last supergame with $N = 4$ is 18.3 percent, but after the reduction to $N = 2$ the cooperation rate immediately jumps up to 60.0 percent. This jump in cooperation as $N$ decreases is then sustained across the remaining supergames, with 58.3 percent cooperation by supergame 20.

Inspecting the results illustrated in Figure 3(A) it is clear that there is minimal evidence for the hypothesis that selected equilibrium is sticky to a within-population shift in $N$. Despite exposure to a prior environment in the first half of the session, longer-run behavior in the second half is not dissimilar from that observed in the between-subject design. This is indicated by the close proximity of the two colored/gray line pairs in supergame 20, and the relative distance from the other pair.[47]

Overall, we find that:

RESULT 3 (Between vs. Within).   *Changing $N$ within subjects as opposed to between does not substantially alter the qualitative results. We find no evidence that the selected equilibrium is sticky in the long run as we shift a primitive within the population.*

---

46.   When testing differences in initial cooperation rates in supergames 6–10 within each $N$ (comparing between and within sessions with identical treatment up to this point), we find $p = 0.150$ for $N = 2$ and $p = 0.981$ for $N = 4$ using $t$-tests. A joint test across both values of $N$ yields $p = 0.353$.

47.   In Online Appendix B, we offer a more detailed like-with-like comparison of the between-subject and within-subject results. These additional findings do not indicate differences with the between-subject results as we move from $N = 2$ to $N = 4$. However, contrary to the hypothesis that the selected equilibrium is sticky, we observe a significant *increase* in responsiveness to changes in $N$ (relative to the between-subject treatments) in the $4 \rightarrow 2$ treatment.

(A) Between vs. within ($X = \$9$)

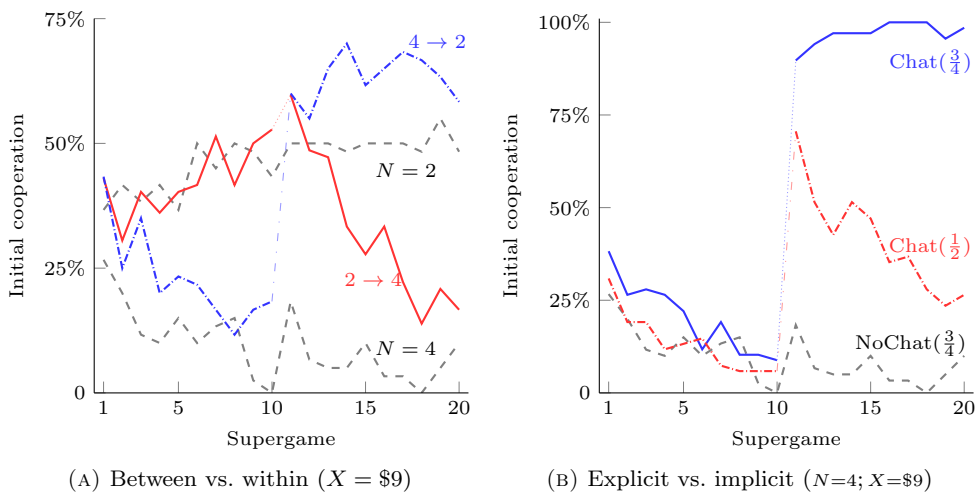(B) Explicit vs. implicit ($N=4$; $X=\$9$)

FIGURE 3. Initial cooperation rates in extensions (by supergame)

## 5.2. Explicit Coordination

In this second set of extension treatments, we examine the strategic-uncertainty mechanism underlying the basin-size model. Specifically, we study the extent to which our results may be influenced by explicit coordination, as free-form communication can diminish strategic uncertainty by enabling players to reveal their strategic intentions.[48] This analysis is motivated by an empirical finding indicating that instances of detected collusion in the industry often originate from explicit collusion—despite the illegality of such meetings.[49]

We design our "chat" treatments by modifying an environment with the least-collusive outcomes, represented by the ($N=4$; $X=\$9$) treatment. In our first chat treatment, Chat(3/4), the initial ten supergames replicate the conditions of the ($N=4$; $X=9$) treatment. However, in supergames 11–20, we introduce pre-supergame chat between all four players. The second chat treatment, Chat(3/4), mirrors Chat(3/4) in terms of timing of when the chat is introduced but reduces the continuation probability to $\delta' = 1/2$ (this continuation is used across all twenty supergames). The Chat(1/2) treatment keeps constant the stage-game payoffs and number of players, but lowers the continuation probability $\delta$ to the point that the grim-trigger strategy is only a knife-edge subgame perfect equilibrium, requiring a critical belief of $p^\star(\delta') = 1$ on

---

48.  Our design is not tailored to pinpoint the exact channel through which strategic uncertainty is reduced. It could be that messages convey the opponent's reasonableness and understanding of the game's tensions. Alternatively, messages might not directly convey information on rationality but simply reduce social distance, making it easier to trust the other player.

49.  See Marshall and Marx (2012) for a more comprehensive treatment.

the other three players cooperating (and so Equations (4) and (5) also coincide). Therefore, Chat(1/2) serves as a litmus test for whether explicit coordination can implement outcomes that are *not* supportable as a robust equilibrium (that is, with arbitrarily small trembles in others' behavior).

In Figure 3(B) we depict initial cooperation rates by supergame, using the ($N$=4; $X$=\$9) treatment as a baseline, here labeled NoChat(3/4).[50] The figure highlights an unambiguous result on the power of explicit coordination under $\delta = 3/4$: providing pre-play chat takes the near-zero initial cooperation rate in NoChat(3/4) to almost full cooperation (98.8 percent, with 80.6 percent ongoing cooperation) in Chat(3/4). Such high levels of cooperation with communication are inconsistent with the predictions of the independent-basin model. Therefore, once explicit coordination devices are allowed for and strategic uncertainty dissipated, our independent basin-size model becomes redundant. That is, the independent model is only intended for implicit/tacit coordination.

However, as we shift the continuation probability to the $\delta' = 1/2$ boundary, even with pre-play communication, participants find it challenging to sustain cooperation. While initial cooperation is substantially higher than the baseline without chat (30.0 percent), ongoing cooperation falls to 4.4 percent (with an ongoing success rate of 0.2 percent). As such, our second chat treatment indicates that for explicit communication to play a role, collusion needs to be at least supportable as non-knife edge equilibrium outcome.

RESULT 4 (Implicit vs. Explicit). *In a multi-player setting, where implicit cooperation results in near-zero cooperation, explicit coordination leads to very high levels of cooperation. However, in the limiting case, where cooperation is a knife-edge subgame perfect equilibrium outcome, even pre-play chat fails to support cooperation.*

### 5.3. Easing Requirements for a Success

In our prior treatments, we find a clear reduction in coordination on the efficient group outcome as we increase $N$. However, in our experiments, as we increase the number of players to four, we are mechanically making it harder to coordinate, as requiring four cooperators out of four is more stringent than requiring two cooperators out of two. In our final robustness exercise, we explore an alternative construction of a four-player game. Specifically, we make it mechanically *easier* to coordinate by allowing for an efficient group-wide outcome even if only two out of four players cooperate. At first sight, relaxing the bar for success in this way suggests that groups will successfully coordinate at much higher rates. However, as we will show, the basin of attraction

---

50. Late-session cooperation and success rates (in supergames 16–20 with subject-clustered standard errors) are provided in Table A.5 in Online Appendix A.

makes the reverse prediction. The reason for this counter-intuitive prediction is that making it mechanically easier to attain a cooperative outcome introduces a new coordination challenge: *who* will cooperate and *who* gets to freeride? In fact, this final extension adds so much strategic uncertainty that our new treatment has a full basin for always defect. As such, this robustness treatment provides a stark test of the basin-of-attraction notion as we increase $N$.

We hold constant the RPD's $2 \times 2$ stage-game representation but weaken the requirement for a success signal to the case where $M - 1$ or more other players cooperate, with $1 \leq M \leq N$. Defining the count of cooperative actions for the other players as $\text{CoopCount}_{N-1}(a_{-i}) := \sum_{j \neq i} \mathbf{1}_{a_j = C}$, an agent's signal is given by:

$$\sigma(a_{-i}; M, N) = \begin{cases} S & \text{if } \text{CoopCount}_{N-1}(a_{-i}) \geq M - 1, \\ F & \text{otherwise,} \end{cases} \tag{6}$$

where for our original treatments $M = N$.

Easing a requirement for success makes it structurally easier to generate a group-wide success. Define $Q_p(M, N)$ as the probability of having $M$ cooperators among $N$ players, where each player chooses to cooperate with probability $p$. For any fixed $p \in (0, 1)$ we have:[51]

$$Q_p(M, N) > Q_p(M, M) > Q_p(N, N). \tag{7}$$

Although it is mechanically easier to achieve joint success for any fixed cooperation rate $p$, weakening the success requirement introduces additional strategic uncertainty. If an individual believes that the other players select a conditionally cooperative strategy with probability $p$, the agent will focus on the following pivotal probability:

$$q_p(M, N) = \Pr\{M - 1 \text{ from } N - 1 \text{ others choose } \alpha_{\text{Grim}}; p\}.$$

In all other situations the agent's action will not affect the long-run outcome: (i) There will either be fewer than $M - 1$ others cooperating (with a miscoordination cost of $(1 - \delta)x$ to the agent); or (ii) There will be $M$ or more cooperators and group-wide success will be guaranteed (with a miscoordination cost of $x$ to the agent for the unnecessary coordination).[52] Therefore, best-responding agents will only cooperate at intermediate values of $p$. The basin of attraction for always-defect will either be full

---

51.    The second inequality is just $p^M > p^N$ which follows as $M < N$. The first inequality comes from decomposing the probability for a group-wide success to the chance the first $M$ players jointly cooperate (meaning it must succeed) and a remaining positive probability, so $Q_P(M, N) = Q_P(M, M) + (1 - Q_P(M, M)) \Pr\{M \text{ cooperate from } N | \text{First } M \text{ not all cooperators}\}$.
52.    The condition for grim to be a best response is:

$$\Pr(\text{Exactly } M - 1 \text{ choose grim}) \geq x \frac{(1 - \delta)}{\delta} + x \Pr(\text{More than } M - 1 \text{ choose grim}).$$

For full derivation see Online Appendix C.

TABLE 4. Basin-effect decomposition: Two-from-four treatment

| Group-wide success | Two-from-four | Compared to: | |
|---|---|---|---|
| | | $(N{=}2; X{=}\$9)$ | $(N{=}4; X{=}\$9)$ |
| All rounds | 0.255 | 0.360 $(p{=}0.006)$ | 0.000 $(p{<}0.001)$ |
| Initial rounds | 0.302 | 0.280 $(p{=}0.751)$ | 0.000 $(p{<}0.001)$ |
| Ongoing rounds | 0.223 | 0.397 $(p{<}0.001)$ | 0.000 $(p{<}0.001)$ |

Results are calculated using data from the last-five supergames. The values in parentheses correspond to $p$-values testing differences between the two-for-four treatment and each of the reference treatments. For the $(N{=}2; X{=}\$9)$ comparisons we use standard tests of proportion; however, because we have no outcome variation in $(N{=}4; X{=}\$9)$, for those tests we use likelihood ratio tests over binomial probabilities.

(all $p \in [0,1]$) or spread across two disjoint regions ($p \in [0, \underline{p}^\star] \cup [\overline{p}^\star, 1]$). A strategic agent will not want to conditionally cooperate: (i) When others cooperate with low probability, $p \in [0, \underline{p}^\star]$, coordination is likely to fail even if they cooperate; or (ii) When others cooperate with high probability, $p \in [\overline{p}^\star, 1]$, coordination is likely to succeed even if they defect. For this reason, if $1 < M < N$, perfectly correlated beliefs result in an always-defect basin size of one for any $x > 0$ and $\delta \in (0,1)$.

Keeping the other parameters constant at $X = \$9$ and $\delta = 3/4$, our final treatment requires $M = 2$ cooperators from the group of $N = 4$ for joint success (which we will call the *two-from-four* treatment). While this change makes successes mechanically easier to obtain than in the $(N{=}2; X{=}\$9)$ and $(N{=}4; X{=}\$9)$ treatments, our basin measure of strategic uncertainty (either independent or correlated) makes the opposite prediction. In fact, at $X = \$9$ and $\delta = 3/4$, the coordination problem is exacerbated to such a degree that weakening the requirement for a success leads to theoretically full basin for always defect.[53,54]

Results from the two-from-four treatment yield a cooperation rate of 22.7 percent across all rounds, and a group-wide success rate of 25.5 percent, compared to a group-wide success rate in the baseline $(N{=}2; X{=}\$9)$ treatment of 36.0 percent. Hence, as predicted by our basin calculations, easing the requirement for success significantly reduces successful coordination $(p = 0.006)$.[55]

---

53.  At $\delta = 3/4$ the always-defect independent-extension basin is smaller than one for $X < X^\star = \$7.91$. For all greater temptations the always-defect basin is full.

54.  We conducted three sessions for the two-from-four treatment (with 64 unique participants). Instructions for this treatment are identical to the four-from-four treatment, except for the explanation of the success/failure signals.

55.  Cooperation at the individual level is also significantly lower in the two-from-four robustness treatment $(p < 0.001)$, compared to the 46.7 percent cooperation rate observed in the two-player RPD. However, as we weaken the cooperation requirement for efficiency, our focus is on a more-comparable measure, successful coordination in the group.

This directional success in a counter-intuitive direction certainly suggests that part of the additional difficulty in coordinating is captured by the basin. However, the basin measure fails to order the success rates for the two-from-four treatment relative to the ($N=4$; $X=\$9$) treatment. While this certainly motivates further research, some caution is warranted. As shown in Figure 2(B) the ongoing cooperation rate in the two-from-four treatment (marked with an empty circles) is not far from what we might expect from the meta-study in RPD games where cooperation is not an equilibrium (basin size of one). In contrast, as we highlighted earlier, the ($N=4$; $X=\$9$) treatment with an ongoing cooperation rate that is almost at zero represents the only treatment that is notably far from the meta-study.

One possible explanation for the result is the stark nature of feedback and learning in the ($N=4$; $X=\$9$) game. With a cooperation rate of 25 percent (approximately the expected value from the meta-study basin), the anticipated group-wide success rate with four players is merely 0.4 percent. In contrast, even with a lower cooperation rate of 20 percent in the two-from-four treatment, we would expect a considerably less extreme group-wide success rate of 18.1 percent.

In summary, we find that:

RESULT 5 (Easing Requirements for a Success).  *In a treatment where the set of players needed for a successful outcome ($M = 2$) is lower than the group size ($N = 4$), the basin-of-attraction extension predicts reduced coordination due to an increase in strategic uncertainty. The treatment results indicate low cooperation rates in line with empirical rates observed for extreme basin-values in other RPD experiments. In terms of successful coordination, the effect from weakening the coordination requirements matches the basin prediction, with a significant decrease in successful coordination relative to the treatment where $M = N = 2$. However, we also find that coordination is higher than in the $M = N = 4$ treatment, which runs counter to the prediction. This finding accentuates the extreme results in our high-tension ($X = 9$) multi-player ($M = N = 4$) game.*

## 6. Conclusion

Our paper examines equilibrium selection in repeated games and the extent to which it can be predicted with a model of strategic uncertainty. We leverage a model of equilibrium selection that rationalizes behavior in the two-player RPD and design an experiment to stress test this specific theoretical model. The predictive model works by mediating the effects from multiple primitives into a single dimension that captures strategic uncertainty. As such, even for rich counterfactual policies with many changes to the setting, the model can still generate a directional prediction. We introduce a novel source of strategic uncertainty that has not yet been studied in the RPD setting (the number of players), while also manipulating a payoff parameter. Therefore, we

can change both sources of strategic uncertainty simultaneously and study the extent to which the evidence is consistent with the predictions of the selection model.

Our main finding is that the model of equilibrium selection can indeed be used as a device to understand successful ongoing coordination on the collusive outcome. In particular, the model performs well in trading off the competing effects from the two distinct sources of strategic uncertainty. Meanwhile, we also document that the model is less successful in predicting initial cooperation rates. Outside of the laboratory, observing the initial round of cooperation can be challenging, but there is more hope that policymakers can observe features of ongoing interactions. Naturally, our game is highly stylized, but it suggests that further research that tests this model of equilibrium selection in more realistic settings may be useful for policy. Given the primitives of an environment, the basin-of-attraction model may be able to predict for what situations ongoing collusion is more likely to emerge. This information might be useful for antitrust authorities to decide which industries to allocate more attention to.

After illustrating the theoretical power of the model for implicit coordination, we turn to several application-motivated extensions that probe the model's limitations. We first show that results continue to hold when manipulations are introduced within the same population. We next document that if subjects are provided with a tool that reduces strategic uncertainty (pre-play chat), the selection model is inappropriate for predicting behavior. That is, the model fails to predict when collusion can be explicitly coordinated. We finally demonstrate that easing cooperation requirements, specifically by stipulating that a subset of players is sufficient to achieve the efficient outcome, does not necessarily promote collusion. This can occur because, with a subset of players being adequate for the efficient outcome, additional strategic uncertainty arises regarding which individuals will cooperate and who may free-ride. The model captures this extra source of strategic uncertainty, predicting decreased cooperation

Taking a step back, a shortcoming of any experimental paper is that conclusions are specific to the chosen environment and parameterizations. Ideally, one would want to evaluate the criterion for equilibrium selection in a large set of repeated games, and in each set for several possible parameterizations. While this goal is outside the scope of the paper, we now outline how we plan to address this in a companion paper (Boczoń, Vespa, and Wilson, 2024) that lays out a possible path for future research in this area.

The idea is that one can evaluate the performance of artificial intelligence algorithms (AIAs) that companies use for pricing decisions (Calvano, Calzolari, Denicolo, and Pastorello, 2020; Asker, Fershtman, and Pakes, 2021) within the RPD setting. The companion paper shows that the experimental results for both the previous RPD literature and our main environments with $N > 2$ can be replicated with AIAs. Given that we find a qualitative and a quantitative match between the long-run behavior of AIAs and our lab participants, the former can be used to predict behavior of

human subjects in counterfactual environments that are not directly studied in the laboratory. Although not as analytically tractable as our basin calculation, such AIAs can be used to expand the scope of experimental studies if partially validated on the narrower domains studied within the laboratory.

# References

Agranov, Marina, Guillaume R Fréchette, Thomas R Palfrey, and Emanuel Vespa (2016), "Static and dynamic underinvestment: An experimental investigation." *Journal of Public Economics*, 143, 125–141.

Asker, John, Chaim Fershtman, and Ariel Pakes (2021), "Artificial intelligence and pricing: The impact of algorithm design." *National Bureau of Economic Research.*

Battaglini, Marco, Salvatore Nunnari, and Thomas R Palfrey (2012), "Legislative bargaining and the dynamics of public investment." *American Political Science Review*, 106, 407–429.

Battaglini, Marco, Salvatore Nunnari, and Thomas R Palfrey (2016), "The dynamic free rider problem: A laboratory study." *American Economic Journal: Microeconomics*, 8, 268–308.

Battalio, Raymond, Larry Samuelson, and John Van Huyck (2001), "Optimization incentives and coordination failure in laboratory stag hunt games." *Econometrica*, 69, 749–764.

Berry, James, Lucas C Coffman, Douglas Hanley, Rania Gihleb, and Alistair J Wilson (2017), "Assessing the rate of replication in economics." *American Economic Review*, 107, 27–31.

Blonski, Matthias and Giancarlo Spagnolo (2015), "Prisoners' other dilemma." *International Journal of Game Theory*, 44, 61–81.

Boczoń, Marta, Emanuel Vespa, and Alistair J Wilson (2024), "Lab to algorithm: Prediting AIs with humans, and vice versa." *In preparation.*

Brandts, Jordi and David J Cooper (2006), "A change would do you good.... an experimental study on how to overcome coordination failure in organizations." *American Economic Review*, 96, 669–693.

Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicolo, and Sergio Pastorello (2020), "Artificial intelligence, algorithmic pricing, and collusion." *American Economic Review*, 110, 3267–3297.

Cason, Timothy N (2008), "Price signaling and 'cheap talk' in laboratory posted offer markets." *Handbook of Experimental Economics Results*, 1, 164–169.

Cason, Timothy N, Tridib Sharma, and Radovan Vadovič (2020), "Correlated beliefs: Predicting outcomes in $2 \times 2$ games." *Games & Economic Behavior*, 122, 256–276.

Cooper, David J and Kai-Uwe Kühn (2014), "Communication, renegotiation, and the scope for collusion." *American Economic Journal: Microeconomics*, 6, 247–78.

Cooper, David J and Roberto A Weber (2020), "Recent advances in experimental coordination games." *Handbook of experimental game theory*, 149–183.

Cooper, Russell, Douglas V DeJong, Robert Forsythe, and Thomas W Ross (1993), "Forward induction in the battle-of-the-sexes games." *American Economic Review*, 1303–1316.

Cooper, Russell, Douglas V DeJong, Robert Forsythe, and Thomas W Ross (1994), "Alternative institutions for resolving coordination problems: experimental evidence on forward induction and preplaycommunication." *Problems of coordination in economic activity*, 129–146.

Cooper, Russell W, Douglas V DeJong, Robert Forsythe, and Thomas W Ross (1990), "Selection criteria in coordination games: Some experimental results." *The American Economic Review*, 80, 218–233.

Crawford, Vincent P, Uri Gneezy, and Yuval Rottenstreich (2008), "The power of focal points is limited: Even minute payoff asymmetry may yield large coordination failures." *American Economic Review*, 98, 1443–1458.

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

Dal Bó, Pedro and Guillaume R Fréchette (2018), "On the determinants of cooperation in infinitely repeated games: A survey." *Journal of Economic Literature*, 56, 60–114.

Dal Bó, Pedro, Guillaume R Fréchette, and Jeongbin Kim (2021), "The determinants of efficient behavior in coordination games." *Games and Economic Behavior*, 130, 352–368.

Davis, Douglas D (2002), "Strategic interactions, market information and predicting the effects of mergers in differentiated product markets." *International Journal of Industrial Organization*, 20, 1277–1312.

Devetag, Giovanna and Andreas Ortmann (2007), "When and why? A critical survey on coordination failure in the laboratory." *Experimental economics*, 10, 331–344.

Dufwenberg, Martin and Uri Gneezy (2000), "Price competition and market concentration: An experimental study." *International Journal of Industrial Organization*, 18, 7–22.

Fonseca, Miguel A and Hans-Theo Normann (2012), "Explicit vs. tacit collusion—the impact of communication in oligopoly experiments." *European Economic Review*, 56, 1759–1772.

Fudenberg, Drew, David G Rand, and Anna Dreber (2012), "Slow to anger and fast to forgive: Cooperation in an uncertain world." *American Economic Review*, 102, 720–749.

Ghidoni, Riccardo and Sigrid Suetens (2022), "The effect of sequentiality on cooperation in repeated games." *American Economic Journal: Microeconomics*, 14, 58–77.

Goette, Lorenz and Armin Schmutzler (2009), "Merger policy: What can we learn from competition policy." *Experiments and Competition Policy; Hinloopen, Jeroen, Hans-Theo Normann, Eds*, 185–216.

Harrington, Joseph E, Roberto Hernan Gonzalez, and Praveen Kujal (2016), "The relative efficacy of price announcements and express communication for collusion: Experimental findings." *Journal of Economic Behavior & Organization*, 128, 251–264.

Harsanyi, John C and Reinhard Selten (1988), *A general theory of equilibrium selection in games.* The MIT Press, Cambridge, MA.

Holm, Håkan J (2000), "Gender-based focal points." *Games and Economic Behavior*, 32, 292–314.

Horstmann, Niklas, Jan Krämer, and Daniel Schnurr (2018), "Number effects and tacit collusion in experimental oligopolies." *Journal of Industrial Economics*, 66, 650–700.

Huck, Steffen, Kai A Konrad, Wieland Müller, and Hans-Theo Normann (2007), "The merger paradox and why aspiration levels let it fail in the laboratory." *Economic Journal*, 117, 1073–1095.

Huck, Steffen, Hans-Theo Normann, and Jörg Oechssler (2004), "Two are few and four are many: Number effects in experimental oligopolies." *Journal of Economic Behavior & Organization*, 53, 435–446.

Kartal, Melis and Wieland Müller (2022), "A new approach to the analysis of cooperation under the shadow of the future: Theory and experimental evidence." Working Paper.

Kartal, Melis, Wieland Müller, and James Tremewan (2021), "Building trust: The costs and benefits of gradualism." *Games & Economic Behavior*, 130, 258–275.

Kloosterman, Andrew (2019), "Cooperation in stochastic games: A prisoner's dilemma experiment." *Experimental Economics*, 23, 447–467.

Lugovskyy, Volodymyr, Daniela Puzzello, Andrea Sorensen, James Walker, and Arlington Williams (2017), "An experimental study of finitely and infinitely repeated linear public goods games." *Games & Economic Behavior*, 102, 286–302.

Marshall, Robert C and Leslie M Marx (2012), *The economics of collusion: Cartels and bidding rings.* The MIT Press, Cambridge, MA.

Martinez-Martinez, Ismael and Hans-Theo Normann (2024), "Cooperation in multiplayer dilemmas." *Working paper*.

Mermer, Ayse Gül, Wieland Mueller, and Sigrid Suetens (2021), "Cooperation in infinitely repeated games of strategic complements and substitutes." *Journal of Economic Behavior & Organization*, 188, 1191–1205.

Potters, Jan and Sigrid Suetens (2013), "Oligopoly experiments in the current millennium." *Journal of Economic Surveys*, 27, 439–460.

Rosokha, Yaroslav and Chen Wei (2020), "Cooperation in queueing systems." *Working Paper.*

Salz, Tobias and Emanuel Vespa (2020), "Estimating dynamic games of oligopolistic competition: An experimental investigation." *RAND Journal of Economics*, 51, 447–469.

Sherstyuk, Katerina, Nori Tarui, and Tatsuyoshi Saijo (2013), "Payment schemes in infinite-horizon experimental games." *Experimental Economics*, 16, 125–153.

Spagnolo, Giancarlo and Matthias Blonski (2001), "Prisoners' other dilemma." *Working Paper.*

Straub, Paul G (1995), "Risk dominance and coordination failures in static games." *The Quarterly Review of Economics and Finance*, 35, 339–363.

Vespa, Emanuel (2020), "An experimental investigation of cooperation in the dynamic common pool game." *International Economic Review*, 61, 417–440.

Vespa, Emanuel and Alistair J Wilson (2019), "Experimenting with the transition rule in dynamic games." *Quantitative Economics*, 10, 1825–1849.

Vespa, Emanuel and Alistair J Wilson (2020), "Experimenting with equilibrium selection in dynamic games." *Working Paper.*

Vesterlund, Lise (2016), "Using experimental methods to understand why and how we give to charity." *Handbook of Experimental Economics*, 2, 91–151.

Weber, Roberto A (2006), "Managing growth to achieve efficient coordination in large groups." *American Economic Review*, 96, 114–126.

Wilson, Alistair J and Emanuel Vespa (2020), "Information transmission under the shadow of the future: An experiment." *American Economic Journal: Microeconomics*, 12, 75–98.