EXPERIMENTING WITH EQUILIBRIUM SELECTION IN DYNAMIC GAMES

EMANUEL VESPA AND ALISTAIR J. WILSON

ABSTRACT. Many economic applications, across an array of fields, use dynamic games to study their questions. Generically, as these games have large sets of possible equilibria, strong selection assumptions are typically required, where Markov perfect equilibrium (MPE) is the preeminent assumption for applied work. Our paper experimentally examines selection across a number of simple dynamic games. Starting from a dynamic modification to the most-studied static environment—the infinitely repeated PD game—we characterize the response to changes to broad qualitative features of a dynamic game. Subjects in our experiments show an affinity for conditional cooperation, adjusting their behavior not only in response to the state but also the recent history. Remarkably, the frequency of MPE-like play in our very different treatments are well organized with an easy-to-compute selection index.

1. INTRODUCTION

The trade-off between opportunistic behavior and cooperation is a central economic tension. In settings where agents interact indefinitely it is theoretically possible to sustain many different outcomes. So long the parties involved place enough weight on the long run and are offered at least their individually rational payoffs, threats to condition future behavior on the present outcome are credible, and powerful enough to deter deviations. This folk-theorem reasoning holds whether the strategic environment is fixed (a repeated game) or evolving through time (a dynamic game), where history-dependent play allows for a large set of subgame-perfect equilibria (SPE). The ensuing indeterminacy in prediction typically requires strong assumptions on equilibrium selection for progress to be made in applied theory and empirical applications.[1] However, these selection assumptions are not innocuous, and can lead to qualitative changes in structural parameter estimates and policy conclusions. While the experimental literature has documented a number of robust patterns for equilibrium selection in repeated games s(see Dal Bó and Fréchette, 2014, for a survey), much less is known for the larger family of dynamic games. Our experiments expand outward from

Date: June, 2017.
We would like to thank Gary Charness, Pedro Dal Bó, John Duffy, Matthew Embrey, Ignacio Esponda, Guillaume Fréchette, Drew Fudenberg, John Kagel, Dan Levin, Tom Palfrey, Ryan Oprea and Lise Vesterlund.

[1]A few examples of dynamic games across a number of fields: Industrial Organization (Maskin and Tirole, 1988; Bajari et al., 2007), Labor Economics (Coles and Mortensen, 2011), Political Economy (Acemoglu and Robinson, 2001), Macroeconomics (Laibson, 1997), Public Finance (Battaglini and Coate, 2007), Environmental Economics (Dutta and Radner, 2006), Economic Growth (Aghion et al., 2001) and Applied Theory (Rubinstein and Wolinsky, 1990; Bergemann and Valimaki, 2003; Hörner and Samuelson, 2009).

these well-understood repeated environments, and we study how behavior and selection respond to qualitative features inherent to dynamic games.

The preeminent solution concept in the dynamic-games literature focuses on the subset of SPE in which agents use Markov strategies (known as Markov perfect equilibria, MPE). Agents using these strategies condition their choices only on the present "state" of the game regardless of the precise history leading to it.[2] Because Markov strategies are memoryless, they cannot punish based on observed deviations from the intended path of play. A consequence of this feature is that Markov strategies cannot support efficient outcomes in many applications of economic interest, where the most-efficient MPE can often be far from the first best. While a key motive for Markov assumptions is analytic tractability, similarly tractable SPEs can be constructed to support an efficient path with a history-dependent trigger that reverts to Markov-play on a deviations. In such SPE, which we refer to as MPE-triggers, players cooperate/collude on first-best choices, with the threat of reversion to the MPE on a deviation. Although the full set of SPE can be intractably large, MPE and MPE-trigger strategies can provide useful benchmarks for understanding selection and our paper provides evidence that behavior is frequently consistent with these two strategies. Though the restriction is useful, the larger question remains: Under which conditions is behavior largely rationalized by history independent play, and when might it be better rationalized by history-dependent play aimed for efficiency?

The class of dynamic games involves a very rich set of environments, and selection can in principle be affected by arbitrary features of the environment. One approach is to implement well-studied models in the laboratory,[3] and the evidence from the nascent experimental literature suggests that behavior in many important dynamic-game environments *is* consistent with the Markov restriction.[4] The strength of the model-by-model approach is that regularities in behavior speak directly to selection in specific environments relevant to applications. A shortcoming is that these experiments are not designed to isolate and test how selection responds to more fundamental features of the dynamic game that are held constant within each model. For instance, it is conceivable that equilibrium selection may be affected by whether the transition between states is exogenous (as

[2]Here we refer to the notion of Markov states, which are endogenously defined as a partition of the space of histories (for details see Maskin and Tirole, 2001). The notion of Markov states is different from the notion of automaton states (for example, a shirk state and a cooperative state in a prisoner's dilemma). For a discussion on the distinction see Mailath and Samuelson (2006, p.178).

[3]For example, Battaglini et al. (2012) study a laboratory version of Battaglini and Coate (2007); Saijo et al. (2014) implement the global warming model of Dutta and Radner (2006); Vespa (2016) implements the dynamic common-pool model of Levhari and Mirman (1980); Salz and Vespa (2016) study the entry-exit framework of Ericson and Pakes (1995). For other experiments with infinite-horizon dynamic games see Rojas (2012), Battaglini et al. (2014), Benchekroun et al. (2014) and Kloosterman (2015).

[4]For example, Battaglini et al. (2012), Battaglini et al. (2014) and Salz and Vespa (2016) provide evidence that the comparative statics are well organized by MPE. In Vespa (2016), the choices of a majority of subjects can be rationalized using Markov strategies.

in Dixit et al., 2000) or endogenous to players actions (for example, Battaglini and Coate, 2007). Likewise, it is possible that the nature of strategic externalities may play a role in selection, where in some applications the agent's current choice only affects others' well-being in the future through the state (as in Levhari and Mirman, 1980), and in other applications agent's choices also affect the others' contemporaneous utility (for example, Acemoglu and Robinson, 2001).

In this paper we follow an alternative but complementary path to the model-by-model approach. We design a family of simple dynamic games around a pivot; a set of laboratory test tubes that turn on and off fundamental building blocks of the dynamic environment. We can therefore identify the qualitative comparative-static effect relative to the pivot: Are subjects reactive to the feature? Does it increase or decrease the selection of MPE-like play?

Our pivot treatment extends the indefinitely repeated PD game by adding a single additional state, so that the stage-game can change across rounds. In both states agents face a PD stage game. However, the payoffs in the *Low* state are unambiguously worse than those in the *High* state. The game starts in *Low* and only if both agents cooperate does it transition to *High.* Once in the *High* state the game transitions back to *Low* only if both agents defect. This game's parameters are chosen so that there is a unique symmetric MPE where agents cooperate in *Low* and defect in *High*, but efficient outcomes that reach and stay in *High* can only be supported with history-dependent play. The modifications to the pivot involve *seven* between-subject treatments, where each examines the effects from a change to an isolated feature of the original game on strategic behavior.

In the pivot, and many of the modifications, we find that a majority of subjects seek to support efficient outcomes with history-dependent play, at comparable levels to those reported for infinitely repeated PD games. This is not to say that Markov play is non-existent in our data, and importantly where we do observe it, it is consistent with the MPE predictions. About one-fifth of the choice sequences in our pivot are consistent with the MPE prediction, while the frequency of non-equilibrium Markov play is negligible.

Our first manipulation set studies the robustness of the pivot results. To do this we shift the efficient frontier in the dynamic game in two different ways, making a single symmetric SPE focal while holding constant the original MPE prediction. A *static* manipulation alters a single payoff at a single state (holding constant the transition rule, we reduce the temptation payoff in the *High* state). A *dynamic* manipulation alters the transition rule between states to make deviations from joint-cooperation relatively less tempting (holding constant the payoffs, we make it harder to remain in *High*). Finally, in a complexity manipulation we perturb the game, adding small-scale exogenous noise to the pivot's payoffs. In all three manipulations the majority of choices remain consistent with history-dependent strategies aimed at efficiency, replicating the pivot finding.

The second set of manipulations focuses on isolating strategic features in dynamic environments, the nature of the externalities. In a standard infinitely repeated PD there is a contemporaneous externality, as each agent's choice directly affects the other's current payoff. But in more-general dynamic environments, an agent's present choice can affect others' payoffs both contemporaneously and in future periods, as an action today can affect the state tomorrow. We refer to the strategic effects from a present choice on others' current payoffs as a *static* externality, and to the effect operating through the transition as a *dynamic* externality.

We remove the pivot's dynamic externality in two distinct ways. In the first, we make the transition between states exogenous, but where both states can be reached (stochastically) along the path. In the second, we remove the dynamics entirely, playing out each of the pivot's stage-games as separate infinitely repeated games. In both manipulations, the only MPE is the stage-game Nash: unconditional joint defection. Relative to the pivot, we observe substantially less cooperation in both treatments. Dynamic externalities are thereby shown to be an important selection factor for the supra-MPE behavior in the pivot. This indicates that the substantial cooperation rates that we find in the pivot are driven by subjects internalizing the consequences of their choices on an endogenously evolving state. Moreover, we document less cooperation in the dynamic game where both states are exogenously reached relative to the separate infinitely repeated PDs. This suggests that the strategic aspects of the transition between states are important, where lower incentives to cooperate in an as-yet unreached state can contaminate and reduce cooperation in states where the myopic incentives to cooperate is higher.

Removing the pivot's static externalities requires that each agent's contemporaneous choice does not affect the other's contemporaneous payoff. We again conduct two separate parameterizations, in which the broad structure of the equilibrium set remains comparable to the pivot: the efficient actions are exactly the same (and involve staying in the *High* state), while the most-efficient MPE still alternates between the *Low* and *High* states. In both parameterizations we find an increase in the frequency of equilibrium Markov play, and a decrease in the frequency of history dependence. The presence of strong static externalities is therefore also identified as an important factor in selection away from the MPE in the pivot game. Where the static externalities between players are weaker, behavior seems to be better rationalized by ignoring past play.

Taken together, our paper's treatments lead to a number of summary conclusions: (i) Having a dynamic strategic environment does not *necessarily* lead to a prevalence of Markov play, and the non-Markov strategies we observe tend to aim for efficiency through trigger strategies. (ii) For those subjects who do use Markov profiles, the MPE is focal. (iii) Supra-MPE selection in the dynamic-PD games are robust to both dynamic and static shifts of the efficient frontier, and a perturbation of the state-space. (iv) The presence of both static and dynamic externalities affects coordination over history-dependent strategies, where removing either strategic externality leads to much greater selection of MPE behavior.

The paper concludes by discussing how our results can contribute to the larger empirical agenda in dynamic games: the development of predictive criteria for equilibrium selection. We use experimental conclusions (i–iv) outlined above in guiding an extension to a unidimensional index derived for the infinitely repeated PD. Remarkably, the index *does* organize the data across our treatments. Moreover, the specific index we develop is computationally tractable and readily applied to generic dynamic games, providing a simple rule-of-thumb calculation for the sensibility of the MPE assumption.

Our work is an initial step, one that suggests constructive directions for further research. First, given that we develop a specific index, it can serve as a basis to design explicit experimental tests of its accuracy and robustness. Second, our findings can guide research into features of dynamic games that we do not explore in this paper. For instance, we focus on very simple two-state environments that allow us to thoroughly document how behavior in a dynamic game differs when "close" to the infinitely repeated PD. Many applications of dynamic games in the laboratory, however, have focused on cases with large state-spaces.[5] The evidence from these games suggest that the selection of Markov strategies may increase with the size of the state-space. Future work can adapt and refine our understanding of whether selection changes as the state space—and the corresponding strategic uncertainty over the more distant future—increases.

2. Experimental Design and Methodology

2.1. **Dynamic Game Framework.** A dynamic game here is defined as n players interacting through their action choices $a_t \in \mathcal{A} := \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$ over a possibly infinite number of periods (indexed by $t = 1, 2, \ldots$). Underlying the game is a payoff-relevant state $\theta_t \in \Theta$ evolving according to a commonly known transition rule $\psi : \mathcal{A} \times \Theta \to \Theta$, so that the state next round is given by $\theta_{t+1} = \psi(a_t, \theta_t)$.[6] The preferences for each player i are represented by a period payoff $u_i : \mathcal{A} \times \Theta \to \mathbb{R}$, dependent on both the chosen action profile a_t and the current state of the game θ_t. Preferences over supergames are represented by the discounted sum (with parameter δ):

$$V_i\left(\{a_t, \theta_t\}_{t=1}^{\infty}\right) = \sum_{t=1}^{\infty} \delta^{t-1} u_i\left(a_t, \theta_t\right). \tag{1}$$

[5] For example, in Battaglini et al. (2012, 2014) and Saijo et al. (2014) the state space is $\mathbb{R}_0^+$, and in Vespa (2016) it is $\mathbb{N}_0$. While many theoretical applications of dynamic games involve large state-spaces there are also well-known cases with few states (for example, Acemoglu and Robinson, 2001)

[6] For simplicity we define a deterministic transition rule here, reflecting the rules in the majority of our environments. More generally though, and for some modifications, we allow the transition rule to be stochastic, with $\psi : \mathcal{A} \times \Theta \to \Delta\Theta$.

Our main set of experiments will examine a number of very simple dynamic environments with an infinite horizon: two players (1 and 2) engage in a symmetric environment with two possible states ($\Theta = \{L(ow), H(igh)\}$) and two available actions, ($\mathcal{A}_i = \{C\text{(ooperate)}, D\text{(efect)}\}$). Any fewer payoff-relevant states, it is an infinitely repeated game. Any fewer players, it is a dynamic decision problem. Any fewer actions, it is uninteresting.

The state in the first period is given by $\theta_1 \in \Theta$ and evolves according to the transition $\psi(\cdot)$. Given a stage game payoff of $u_i(a, \theta)$ for player i, symmetry of the game enforces $u_1((a, a'), \theta) = u_2((a', a), \theta)$ for all $(a, a') \in \mathcal{A} := \mathcal{A}_1 \times \mathcal{A}_2$ and all states $\theta \in \Theta$.

2.2. **Treatments.** A treatment will be pinned down by the tuple $\Gamma =< \Theta, \theta_1, u_i, \psi >$ indicating a set of possible states Θ, a starting state θ_1, the stage-game payoffs $u_i(a_t, \theta_t)$, and the transition rule $\psi(a_t, \theta_t)$. All other components (the set of actions $\mathcal{A}$ and the discount parameter δ) will be common. In terms of organization, sections 3–5 will describe treatments and results sequentially. After specifying and motivating each treatment, we provide more specific details with respect to the theoretical predictions within each section. In particular, for each treatment we will focus on characterizing symmetric Markov perfect equilibria (MPE, formally defined in the next section) and providing examples of other SPE that can achieve efficient outcomes by conditioning on the history of play.

2.3. **Implementation of the infinite time horizon and session details.** Before presenting treatments and results, we first briefly note the main features of our experimental implementation. To implement an indefinite horizon, we use a modification to a block design (cf. Fréchette and Yuksel 2013) that guarantees data collection for at least five periods within each supergame. The method, which implements $\delta = 0.75$, works as follows: At the end of every period, a fair 100-sided die is rolled, the result indicated by Z_t. The first period T for which the number $Z_T > 75$ is the final payment period in the supergame.

However, subjects are not informed of the outcomes Z_1 to Z_5 until the end of period five. If all of the drawn values are less than or equal to 75, the game continues into period six. If any one of the drawn values is greater than 75, then the subjects' payment for the supergame is the sum of their period payoffs up to the first period T where Z_T exceeds 75. In any period $t \geq 6$, the value Z_t is revealed to subjects directly after the decisions have been made for period t.[7] This method

[7]This design is therefore a modification of the block design in Fréchette and Yuksel (2013), in which subjects learn the outcomes Z_t once the block of periods (five in our case) is over. We modify the method and use just one block plus random termination in order to balance two competing forces. On the one hand we would like to observe longer interactions, with a reasonable chance of several transitions between states. On the other, we would like to observe more supergames within a fixed amount of time. Our design helps balance these two forces by guaranteeing at least five choices within each supergame (each supergame is expected to have 5.95 choices). Fréchette and Yuksel (2013) show that "block designs" like ours can lead to changes in behavior around the period when the information on $\{Z_t\}_{t=1}^{5}$ is

implements the expected payoffs in (1) under risk neutrality. For payment, we randomly select four of the fifteen supergames.[8]

All subjects were recruited from the undergraduate student population at the University of California, Santa Barbara. After providing informed consent, they were given written and verbal instructions on the task and payoffs.[9] Each session consists of 14 subjects, randomly and anonymously matched together across 15 supergames. We conducted at least three sessions per treatment, where each session lasted between 70 and 90 minutes, and participants received average payments of \$19.[10]

2.4. **Overview of the design.** In total we will document results from eight distinct treatments, across two broad categories of manipulation: (i) robustness of our pivot game (Section 4); and (ii) changing strategic externalities, how one agent's choice affects the other's payoffs (Section 5). In each manipulation we change a single feature of our pivot, endeavoring to hold other elements constant. Though we will provide more specific details as we introduce each treatment, the reader can keep track of the full design and the differences across treatments by consulting Table 1.

Table 1 summarizes the main differences for each treatment, relative to the pivot. The table provides: (i) the size of the state-space; (ii) the action profile/probability of transition to a different state; (iii) the starting state θ_1; (iv) the most-efficient symmetric MPE; (v) the efficient action profile; and (vi) the action that obtains the individually rational payoff (by state). However, rather than presenting the entire global design all at once, we introduce each manipulation and its results, in turn. The natural point to begin then is describing our pivot treatment and outlining the behavior we find within it, which we do in the next section. Table 7 in the appendix provides a summary of the central theoretical predictions by treatment.

revealed. However, such changes in behavior tend to disappear with experience and they show that this does not affect comparative statistics across treatments.

[8]Sherstyuk et al. (2013) compare alternative payment schemes in infinitely repeated games in the laboratory. Under a 'cumulative' payment scheme similar to ours subjects are paid for choices in all periods of every repetition, while under the 'last period' payment scheme subjects are paid only for the last period of each supergame. While the latter is applicable under any attitudes towards risk, the former requires risk neutrality. However, Sherstyuk et al. observe no significant difference in behavior conditional on chosen payment scheme, concluding that it "suggests that risk aversion does not play a significant role in simple indefinitely repeated experimental games that are repeated many times."

[9]Instructions are provided in Online Appendix C. In the instructions we refer to periods as rounds and to supergames as cycles.

[10]One treatment has four sessions (En-DPD-CC with 56 subjects), where all others have three sessions (42 subjects).

TABLE 1. Treatment Summary

Treatment	$\|\Theta\|$	**Transition**		$\Pr\{\theta_1 = L\}$	**MPE**		**Efficient**		**IR action**	
		L	H		L	H	L	H	L	H
	Pivot (Section 3):									
En-DPD	2	(C,C)	(D,D)	1	C	D	(C,C)	(C,D)	D	C
	Robustness of the Pivot (Section 4):									
En-DPD-CC	=	=	=	=	=	=	=	(C,C)	=	=
En-DPD-HT	=	=	not (C,C)	=	=	=	=	(C,C)	D	D
En-DPD-X	22	$\doteqdot$	$\doteqdot$	$\doteqdot$	$\doteqdot$	$\doteqdot$	$\doteqdot$	$\doteqdot$	=	=
	Change Strategic Externalities (Section 5):									
Ex-DPD	=	prob. 0.6	prob. 0.2	=	D	D	=	=	D	D
Ex-SPD	1	$\emptyset$	$\emptyset$	prob. 0.4	D	D	=	=	D	D
En-DCP-M	=	=	=	=	=	=	=	=	=	=
En-DCP-E	=	=	=	=	=	=	=	=	D	D

Note: Where the table lists "=", the relevant cell is identical to the En-DPD game's value. For the En-DPD-X we list $\doteqdot$ to indicate similarity on the path, given a changed state-space. The *Transition* column indicates either the action profile a that changes the state (so that $\psi(a,\theta) \neq \theta$) for deterministic transitions or the exogenous probability the state changes given a random transition.

TABLE 2. En-DPD

θ=Low

		2: **C**	2: **D**
1:	**C**	100,100	30, 125
	D	125, 30	60,60

θ=High

		2: **C**	2: **D**
1:	**C**	200, 200	130, 280
	D	280, 130	190, 190

3. PIVOT TREATMENT

3.1. **Pivot Design (En-DPD).** Our pivot game uses two PD stage games, one for each state, and so we label it a dynamic prisoner's dilemma (*DPD*). The transition between the two states is endogenous (*En-*), with a deterministic relationship to the current state and player actions. We therefore label the pivot treatment as "En-DPD."

The precise stage-game payoffs $u_i(a,\theta)$ are given in Table 2 in US cents. The game starts in the low state ($\theta_1 = L$), and the next period's state $\theta_{t+1} = \psi(a_t, \theta_t)$ is determined by

$$\psi(a,\theta) = \begin{cases} H & \text{if } (a,\theta) = ((C,C), L) \\ L & \text{if } (a,\theta) = ((D,D), H) \\ \theta & \text{otherwise.} \end{cases}$$

This transition rule has a simple intuition: joint cooperation in the low state is required to shift the game to the high state; once there, so long as both players don't defect, the state remains in high. Examining the payoffs in each state, both stage games are clearly PD games: D is a dominant strategy but (D, D) is not efficient. Each stage game therefore has a static strategic externality, where the choice of player i in period t alters the period payoff for player $j \neq i$. However, because the transition between states depends on the players' actions the game also has a dynamic externality. The choice of player i in period t affects future states and thus has a direct implication for the continuation value of player j.

An economic interpretation for the pivot game is that the state represents the stock of a good (fish in a pond, water in reservoir, the negative of pollution levels, market demand) and the actions are a choice over that stock (extraction of fish or water, effluent from production, market supply). By cooperating in the low state, the current low stock can be built up to a socially-desirable level. Once the stock is built up to the high level, all actions become more profitable. In conjunction with this, the stock becomes more robust, and only transitions back to low following more systemic opportunistic behavior (joint defection). However, though the payoffs from cooperative behavior increase in the high state, so do the relative temptations for defections, and the relative loss from cooperating when the other defects.

Theoretical Properties. Much of our paper will focus on symmetric Markov strategy profiles, a function $\sigma : \Theta \to \mathcal{A}_i$. Markov strategies only condition on the current state θ_t, ignoring other components of the game's history $h_t = \{(a_s, \theta_s)\}_{s=1}^{t-1}$, in particular the previously chosen actions. Given just two states, there are four possible pure Markov strategies available to each player in our pivot game, an action choice $\sigma_L \in \{C, D\}$ for the low state, and $\sigma_H \in \{C, D\}$ for the high state. We will use the notation $M_{\sigma_L \sigma_H}$ to refer to the Markov strategy

$$\sigma(\theta) = \begin{cases} \sigma_L & \text{if } \theta = L, \\ \sigma_H & \text{if } \theta = H. \end{cases}$$

A symmetric pure-strategy Markov perfect equilibrium (MPE) is a profile $(M_{\sigma_L \sigma_H}, M_{\sigma_L \sigma_H})$ that is also an SPE of the game. For our pivot there is a unique symmetric MPE, the strategy M_{CD}: both players cooperate in low, both defect in high. As such, the path of play for this MPE cycles between the low and high states forever, and the discounted-average payoff is $4/7 \cdot 100 + 3/7 \cdot 190 \simeq 138.6$.

Symmetric profiles that cooperate in the high state, either M_{CC} or M_{DC}, are not sub-game perfect. A single player deviating in the high state increases their round payoff to 280 from 200, but the deviation affects neither the state nor action choices in future periods, so the continuation value is unchanged and the deviation is beneficial. Moreover, the strategy M_{DD} that plays the Nash stage-game in both states is also not an SPE. For any sub-game where the game is in high, this Markov profile dictates that both agents jointly defect from this point onward, yielding the discounted-average payoff $\frac{1}{4} \cdot 190 + \frac{3}{4} \cdot 60 = 92.5$. But the individually rational (IR) payoff in the high state is 130, which each player can guarantee by cooperating in every period. So M_{DD} is not an MPE.[11]

From the point of view of identifying Markov behavior, our parameterization of the pivot game was chosen so that the equilibrium strategy M_{CD} has the following properties: (i) the MPE path transits through both states; (ii) the strategy requires both that subjects do not condition on the history, but also that they select *different* actions in *different* states, and is therefore more demanding than an unconditional choice (for instance, M_{DD}); and (iii) more-efficient SPE are possible when we consider strategies that can condition on history, as we discuss next.

Keeping the game in the high state is clearly socially efficient—payoffs for each player i satisfy $\min_a u_i(a, H) > \max_a u_i(a, L)$. Joint cooperation in both states is one outcome with higher payoffs than the equilibrium MPE, achieving a discounted average payoff of 175. One simple form of history-dependent strategy that can support this outcome in a symmetric SPE is a trigger. Players cooperate in both states up until they observe an action profile $a_{t-1} \neq (C, C)$, after which the trigger is pulled and they switch to an incentive-compatible punishment. One way to make sure the punishment is incentive compatible is to simply revert to the MPE strategy M_{CD}. We will refer to this symmetric history-dependent trigger strategy with an M_{CD} punishment phase as S_{CD}.[12]

Though joint-cooperation is more efficient than the MPE, it is possible to achieve greater efficiency still. The efficient path involves selecting C in the first period and any sequence of actions $\{a_t\}_{t=2}^{\infty}$ such that each $a_t \in \{(C, D), (D, C)\}$. From period two onwards, efficient outcomes yield a total period payoff for the two players of 410, where joint cooperation forever yields 400. Notice, however, that the payoff difference between the efficient solution and joint cooperation is small, amounting to 5 cents per player.[13] One simple asymmetric outcome involves alternating forever between (C, D)/(D, C) in odd/even periods once the game enters the high state. Such an outcome can be supported with an M_{CD}-trigger after any deviation from the intended path, where we will

[11] Expanding to asymmetric MPE, there is an equilibrium where one agent uses M_{DC} and the other M_{DD}. If the starting state were high, this asymmetric MPE can implement an efficient outcome where one agent selects C, the other D, and thereby remain in high. However, since the initial state is low, this strategy will never move the game to the high state, and as such implements the highly inefficient joint-defection in low forever outcome.

[12] The symmetric profile (S_{CD}, S_{CD}) is an SPE for all values of $\delta \geq 0.623$, and so constitutes a symmetric SPE for our pivot game at $\delta = 0.75$. Trigger-strategies where both players punish using M_{DD} (which we call S_{DD}) are not sub-game perfect. However, jointly-cooperative outcomes can be sustained using an asymmetric Markov trigger. In this modification, the player who deviates switches to M_{DC}, while the player who was deviated upon switches to M_{DD}. That is, this strategy uses the asymmetric MPE described in footnote 11 and implements a punishment path of permanent defection. This strategy is a symmetric SPE for all values of $\delta \geq 0.534$ (note that symmetry in action is broken by the observed history, and so both players using this strategy is a symmetric SPE).

[13] We parametrize our pivot treatment with an asymmetric efficient outcome as this baseline leads to a clearer comparisons (similar efficient frontiers) when examining the manipulations of the strategic externalities in Section 5. We provide a detailed explanation of this choice in footnote 38 of Section 5.2. Section 4 will present two treatments where symmetry is efficient, and demonstrates that this feature of the game is not driving our results.

subsequently refer to this asymmetric trigger strategy as A_{CD}. The discounted-average payoff pair from the first period onwards is $(170.7, 186.8)$, so the player who cooperates first suffers a loss relative to joint-cooperation forever. Moreover, though fully efficient outcomes are not attainable with symmetric SPE, or through any type of MPE, *every* efficient outcome in En-DPD is supportable as an SPE for $\delta = 0.75$.[14,15]

3.2. **Pivot Results.** All results in all treatments in this paper are summarized by two figures and a table positioned at the end of this paper.[16] The two figures are designed to illustrate aggregate-level behavior (Figure 2) and variation across supergames (Figure 3), while the table (Table 4) provides estimates of the selection frequency for a number of key strategies. While more-detailed regressions are included in the paper's appendices, to simplify the paper's exposition we will focus on just these three main sources to discuss our results, with details in footnotes and the appendix.

The first source, Figure 2, presents the most-aggregated information on behavior, the average cooperation rate by state, as well as some basic patterns for behavior within and across supergames. The leftmost six bars present results for the En-DPD treatment. The first three gray bars indicate the cooperation rate when the state is low, where the first, second and third bars show averages for supergames 1–5, 6–10 and 11–15, respectively. The height of the bars indicate that the overall cooperation rate in low is close to 75 percent, and is relatively constant as the sessions proceed (albeit with a slight decrease in the last five supergames).

Similarly, the three white bars present the average cooperation rates for all periods in the high state, again across each block of five supergames. The figure illustrates an average cooperation rate in the high state of just under 70 percent in the first five supergames, falling to a little over 50 percent in the last five supergames. These raw numbers suggest that a majority of choices are more-efficient than the MPE prediction of no cooperation in high. More than this though, our data also suggests that at least some subjects are not conditioning solely on the state, that the frequency of cooperation at each state falls as the supergame proceeds. To illustrate this, Figure 2 displays cooperation rates in the first (second) period of each supergame conditional on being in the low (high) state with gray (white) circles. For comparison, the arrows on each bar point to the cooperation rate in the last two periods in each supergame (again, conditional on the relevant state). For En-DPD, the illustrated

[14]Efficient paths must have both players cooperate with probability one in the initial low state and have zero probability of either joint defection or joint cooperation in high. This rules out symmetric mixtures without a correlation device (effectively putting a non-payoff relevant variable into the state-space).

[15]Every attainable point on the efficient frontier for En-DPD is supportable as an SPE for $\delta \geq 0.462$. The bound on δ comes from the one-period deviation in period two and onwards for the following strategy: In period one, both agents cooperate. In period two and beyond, one agent plays C, the other D, with a triggered (M_{DC}, M_{DD}) punishment if the game is ever in the low state in period 2 onward. All other efficient outcomes weaken the temptation to deviate.

[16]As we introduce treatments we will refer back to these three tables frequently. Readers are advised to either bookmark the pages that contain them, or print out additional copies. More-detailed tables with formal statistical tests, the most-common sequences of the state and action choices within supergames are given in the Online Appendix.

pattern shows much higher initial cooperation levels in the low state, approaching 100 percent in the last five supergames. However, the low-state cooperation rate near the end of the supergame is much lower, closer to 50 percent.[17]

To further disaggregate behavior we move to Figure 3, where the unit of observation is the sequence of choices made by each subject in each supergame, which we will refer to as a history. Each history is represented as a point: a cooperation rate in the low state (horizontal axis), and in the high state (vertical axis). The figure rounds these cooperation rates to the nearest tenth (and so the figure can be thought of as an 11×11 histogram) illustrating the number of observed pairs at each point with a bigger circle to represent a greater mass of observations.[18]

Figure 3(A) shows that while most histories in the pivot present a perfect or near-perfect cooperation rate in the low state, the dispersion is much larger along the vertical axis, suggesting the presence of three broad categories of cooperation in the high state. The mass of histories near the top-right corner represent supergames where the choices come close to perfectly cooperative, as predicted by the symmetric history-dependent S_{CD} strategy. The mass in the bottom-right corner has very low high-state cooperation rates, and is consistent with the MPE strategy M_{CD}. Finally, there is a group with high-state cooperation rates close to 50 percent, which could be consistent with the asymmetric A_{CD} strategy that alternates between C and D in the high state to achieve an efficient outcome. However, other strategy pairs might also produce these patterns.

To further inquire which strategies best represent the subjects' choices we use a strategy frequency estimation method (SFEM, for additional details see Dal Bó and Fréchette, 2011).[19] The method considers a fixed set of strategies, and compares the choices that would have been observed had the subject followed the strategy perfectly (taking as given the other player's *observed* actions). Using an independent probability $1 - \beta$ of making errors relative to the given strategy, the process measures the likelihood that the observed choice sequence was produced by each strategy. The method then uses maximum likelihood to estimate a mixture model over the specified strategy set (frequencies of use for each strategy) as well as a goodness-of-fit measure β, the probability any choice in the data is predicted correctly by the estimated strategy mixture.[20]

[17] Table 8 in the appendix provides the cooperation levels by state obtained from a random-effect estimate, while Table 10 explicitly tests whether the *initial* cooperation rate in each state is different than in subsequent periods.

[18] When a history never reaches the high state it is not possible to compute the cooperation rate in high. Such cases are represented in the vertical axis with 'NaN' for not a number.

[19] SFEM has also been used in many other papers, in particular Fudenberg et al. (2010), who also conduct a Monte-Carlo exercise to validate the procedures consistency.

[20] Although the set of included strategies is simple, our measures of goodness-of-fit across treatments are far from a random draw (a β value of 0.5). This suggests that even with this limited set of strategies it is possible to rationalize the data fairly well. Moreover, in Table 16 in the Appendix we reduce the set of strategies even further, to just the MPE and the MPE trigger, M_{CD} and S_{CD}, where the fit parameter is 0.72.

For the estimations reported in Table 4 we specify a very simple set of strategies. It includes all four pure Markov profiles, M_{CC}, M_{DD}, M_{CD} and M_{DC}. In addition, the estimation procedure also includes four strategies that aim to implement joint cooperation. First, we include the two symmetric trigger strategies, S_{CD} and S_{DD}, which differ in the severity of their punishments. We also include two versions of tit-for-tat (TfT). The standard version starts by selecting C in period one and from the next period onwards selects the other's previous-period choice, where this strategy has been documented as a popular choice in previous infinitely repeated PD studies despite not being sub-game perfect. The only difference in the suspicious version ($STfT$) is that it starts by defecting in the first period. We also include two history-dependent asymmetric strategies that seek to implement an efficient, alternating outcome: A_{CD} and A_{DD}, where the difference between the two is again on the triggered punishment after a deviation.[21]

The SFEM estimates for the pivot treatment, available in the first column of Table 4, reflect the heterogeneity observed in Figure 3(A). A large mass of behavior is captured by three statistically significant strategies with comparable magnitudes: M_{CD}, S_{CD} and TfT. The frequency of the MPE strategy is slightly higher than one-fifth and reversion to that strategy is the most popular among those using triggers to achieve joint cooperation, where these trigger strategies (S_{CD} and S_{DD}) capture approximately 30 percent of the estimates.[22]

The mass attributed to TfT represents approximately one-quarter of the estimates. In the En-DPD game, though TfT is a not a symmetric Nash equilibrium, the strategy does provide substantial flexibility. If paired with another subject using TfT, the outcome path results in joint cooperation. However, when paired with other players that defect the first time the high-state is reached TfT can produce an efficient path, and can be part of a Nash equilibrium (in particular, when paired to A_{CD} or A_{DD} which lead with defection in high). TfT is therefore capable of producing both joint cooperation and efficient alternation across actions in the high-state depending on the behavior it is matched to.

3.3. **Conclusion.** The majority of the data in our pivot is inconsistent with the symmetric MPE prediction of joint cooperation in low and joint defection in high. Though we do find that close to

[21]Efficient asymmetric SPE not only require coordination over the off-the-path punishments to support the outcome, they also require coordination over breaking symmetry the first time play reaches high. The strategy specifies that one agent starts by selecting C, and the other D the first time the high state is reached. From then both play the action chosen by the other player last period so long as the outcome is not (D, D), switching to the punishment path otherwise. The appendices present the SFEM output with both strategy sub-components $A_{CD} = \left(A^C_{CD}, A^D_{CD}\right)$ and $A_{DD} = \left(A^C_{DD}, A^D_{DD}\right)$, where A^a_X is the strategy which starts with action a the first time the game enters the high state (see Table 15) and reverts to M_X on any deviation. However, because the two versions of each strategy only differ over the action in one period it is difficult for the estimation procedure to separately identify one from the other. For simplicity of exposition Table 4 includes only the version in which the subject selects D in the first period of the high state, A^D_{CD} and A^D_{DD}.

[22]Outside of the SFEM estimates, in Table 11 in the Appendix we provide direct evidence for history dependence, where continued cooperation in the high state is significantly reduced if the partner defects.

one fifth of subjects are well matched by the M_{CD} strategy profile, many more attempt and attain efficient outcomes that remain in the high state. Over 60 percent of the estimated strategies are those that when matched with one another keep the game in the high state forever through joint cooperation (M_{CC}, S_{DD}, S_{CD} and TfT).

Looking to strategies detected in the infinitely repeated PD literature provides a useful benchmark for comparison. Dal Bó and Fréchette (2014) find that just three strategies account for the majority of PD game data—*Always defect*, the *Grim trigger* and *Tit-fot-Tat*. Through the lens of a dynamic game, the first two strategies can be thought of as the MPE and joint-cooperation with an MPE trigger. The strategies used in our dynamic PD game therefore mirror the static PD literature, where three strategies account for over 60 percent of our data: the MPE M_{CD}; joint cooperation with an MPE trigger, S_{CD}; and tit-for-tat.

Despite the possibility for outcomes with payoffs beneath the symmetric MPE (in particular through the myopic strategy M_{DD} which defects in both states) the vast majority of outcomes and strategies are at or above this level, even where history-dependent punishments are triggered. The MPE strategy is clearly a force within the data, with approximately 40 percent of the estimated strategies using it directly or reverting to it on miscoordination. However, the broader results point to history-dependent play as the norm. The next two sections examine how modifications to the strategic environment alter this finding.

4. Robustness of the Pivot

Our pivot requires several specific choices for the game, through the payoffs and the transition rule. As a robustness check, we now examine a series of three modifications to the pivot. Our first two modifications examine changes to the efficient action. The first reduces the static temptation to defect in the high state, holding constant the continuation value from a defection. The second reduces the continuation value from a defection holding constant the static temptation. Our final robustness treatment adds small non-persistent observable shocks to the game's payoffs. The effect is a substantial perturbation of the game's state space, but without fundamentally altering the strategic tensions.

4.1. **Efficient Actions (Static Efficiency, En-DPD-CC).** Our first modification shifts the efficient actions by decreasing the payoff $u_i\left(\left(D,C\right),H\right)$ from 280 to 250. All other features of the pivot—the starting state, the transition rule, all other payoffs—are the same. The change therefore holds constant the MPE prediction (cooperate in low, defect in high) but reduces the payoffs obtainable with combinations of (C,D) and (D,C) in high. Where in En-DPD the asymmetric outcomes produce a total payoff for the two players of 410, in the modification it is just 380. Joint cooperation in high is held constant, so that the sum of payoffs is 400, as in the pivot. The history-dependent trigger S_{CD} is still a symmetric SPE of the game, but its outcome is now first best, and the temptation to deviate from it is lowered. As the main change in the game is to make the high-state action (C, C) more focal, we label this version of our endogenous-transition DPD game: *En-DPD-CC*.

The data, presented in Figures 2 and 3(B), displays many similar patterns (and some important differences) with respect to the pivot. Initial cooperation rates in both states and both treatments start out at similar levels, but the pattern of declining high-state cooperation across the session observed in En-DPD is not mirrored in En-DPD-CC. High-state cooperation rates for the two treatments are significantly different (at 90 percent confidence), but only for the last five supergames.[23] Looking at the supergame level in Figure 3(B), this increase is reflected through greater concentrations in the top-right corner, perfectly cooperative supergames.

The estimated strategy weights in Table 4 indicate higher frequencies for strategies aimed at joint cooperation. Strategies that lead to joint cooperation when matched (S_{CD}, S_{DD}, TfT and M_{CC}) amount to 70 percent of the estimated frequencies, an increase of ten percentage points over the pivot. The estimated frequency of MPE play is diminished substantially, both directly as the M_{CD} strategy is not statistically significant, and indirectly through miscoordinations, as the symmetric trigger with the most weight is the harsher trigger S_{DD}.

Like the En-DPD results, the large majority of outcomes in En-DPD-CC intend to implement more-efficient outcomes than the MPE. The manipulation in En-DPD-CC makes joint cooperation focal and so easier to coordinate on, and our data reflects this with an even weaker match to the MPE than in the pivot. Our next treatment examines a similar exercise where we instead weaken the continuation value on a defection from joint-cooperation.

4.2. **Transition Rule (Dynamic Efficiency, En-DPD-HT).** In the previous two treatments we discussed, once the game reaches the high state, only joint defection moves it back to low. While the last treatment modified a pivot stage-game payoff so that joint cooperation is first best, our next treatment accomplishes the same thing through a change to the transition rule. Exactly retaining the stage-game payoffs from En-DPD (Table 2) we alter the transition rule in the high-state $\psi(a, H)$ so that any action *except* joint-cooperation switches the state to low next period. The complete transition rule for the state is therefore

$$\theta_{t+1} = \psi(a_t, \theta_t) = \begin{cases} H & \text{if } a_t = (C, C), \\ L & \text{otherwise.} \end{cases}$$

As we are changing the high-state transition (HT) rule, we label the treatment *En-DPD-HT*.

[23] Statistical tests are reported in the appendix's Table 9 from a random-effects probit clustering standard errors at the session level.

There are two broad changes relative to En-DPD: (i) the efficient action in the high state becomes (C, C), as any defection yields an inefficient switch to low next period; and (ii) the individually rational payoff in high is reduced. In the pivot, conditional on reaching the high state, each player can ensure themselves a payoff of at least 130 in every subsequent period by cooperating. However, in En-DPD-HT no agent can unilaterally keep the state in high, as doing so here requires *joint* cooperation. The individually rational payoff in the high state therefore shrinks to $1/4 \cdot 190 + 3/4 \cdot 60 = 92.5$, with the policy that attains the minmax shifting to M_{DD} (where it is M_{DC} in the pivot).

The most-efficient MPE of the game starting from the low state is the same as the pivot (M_{CD}), where the sequence of states and payoffs it generates is identical to that in En-DPD. However, the change in transition rule means that both M_{DD} and M_{DC} are now also symmetric MPE, though with lower payoffs than M_{CD}.[24] Efficient joint cooperation is attainable as an SPE with either symmetric trigger, S_{DD} and S_{CD}.[25]

On the one hand, this change in the transition rule makes supporting an efficient outcome easier. First, joint cooperation is focal, which may aid coordination. Second, the transition-rule change reduces the temptation in the high state, any deviation leads to low for sure next period, and so is less appealing. On the other hand, the changed transition rule may also increase equilibrium Markov play. In En-DPD an agent deviating from M_{CD} in the high state suffers a static loss (a 130 payoff versus 190) that is partially compensated with an increased continuation (next period the game will still be in high). However, in En-DPD-HT there is no reason at all to deviate from M_{CD} in the high state. A one-shot deviation produces both a realized static loss and no future benefit either from a different state next period. For this reason, coordinating away from the MPE strategy M_{CD} becomes harder in En-DPD-HT.

While ex-ante the change in the transition rule could plausibly lead to either more or less MPE play, the data displays a substantial increase in the selection of efficient outcomes. Looking at the state-conditioned cooperation rates in Figure 2 and comparing En-DPD-HT to the pivot, the most apparent results are the significant increase in high-state cooperation.[26] Comparing Figures 3(A) and (C) shows a clear upward shift, with the vast majority of histories in the upper-right corner, tracking instances of sustained joint cooperation. Finally, the SFEM output in Table 4 indicates a substantial increase in strategies involving joint cooperation along the path: adding M_{CC}, S_{DD} and TfT, the total frequency is 91.2 percent.

[24]If the dynamic game were to begin in the high state, the MPE M_{DC} yields an efficient outcome, as it effectively threatens a reversion to the worst-case MPE path if either player deviates. However, given that our game sets $\theta_1 = L$, the path of play for this strategy is inefficient, as it traps the game in low forever.

[25]TfT is a Nash equilibrium of the game, but not an SPE, as there is a profitable one-shot deviation along paths that deviate from joint cooperation.

[26]The difference is significant at the 99 percent confidence level for the last five supergames.

While there is a clear increase in play that supports the efficient symmetric outcome, the SFEM estimates also indicate a shift for the most-popular punishments. In the pivot (and En-DPD-CC) the most-popular history-dependent strategy is TfT. But in En-DPD-HT the most-popular strategy corresponds to the harshest individually rational punishment: S_{DD}, the grim trigger.

We find no evidence for the most-efficient MPE, either directly through M_{CD}, or through subjects using it as a punishment on miscoordination with S_{CD}. The only Markov strategy with a significant estimate is M_{CC}, which is harder to separately identify from history-dependent strategies that succeed at implementing joint cooperation, and is the *only* Markov strategy *inconsistent* with some MPE.[27]

4.3. **State Space Perturbation (En-DPD-X).** One possible reason for the failure of the MPE predictions in our pivot is that the state-space is too simple. History-dependent strategies are common in experiments on the infinitely repeated PD games, with just one state. At the other extreme with an infinite number of states there is experimental evidence for Markov play (cf. Battaglini et al., 2014; Vespa, 2016). One potential selection argument for state-dependent strategies is simply the size of the state-space, where the 20 percent MPE play we observe in our pivot would, ceteris paribus, increase as we add more state variables. Our final robustness treatment examines this idea by perturbing the pivot game in a way that increases the size of the state-space. In so doing, we assess whether the presence of a richer state-space leads to a greater frequency of cognitively simpler Markov strategies.

One simple way to add states while holding constant many of the pivot's strategic tensions is payoff-relevant noise. This treatment adds an commonly known iid payoff shock each period through a uniform draw x_t over the support $X = \{-5, -4, \ldots, 4, 5\}$.[28] The specific payoffs in

[27]The SFEM can identify two strategies that implement joint cooperation only if we observe enough behavior in the punishment phases to separately identify the strategies; otherwise, strategies such as S_{DD}, S_{CD} and M_{CC} are identical. This issue combines with out coarse set of specified strategies. When the procedure reports an estimate for M_{CC}, it can be capturing either unconditional cooperation, M_{CC}, or any history-dependent strategy that mostly cooperates and either does not enter its punishment phase within our data, or where the punishment path is closer to M_{CC} than our other coarsely specified strategies. In particular, any forgiving strategies which take a number of deviations to enter the punishment are likely to be reported as M_{CC} in our estimates. Vespa (2016) develops an experimental procedure to obtain extra information that allows to distinguish between such strategies and gain more identifying power.

[28]Clearly, another test of how changes in the state-space affect behavior would involve increasing the size of states that are reached endogenously. A thorough study of such manipulations is outside of the scope of this paper. The exogenous shocks that we study can be thought of as a creating a small perturbation of the pivot game. From another point of view, exogenous shocks are common in IO applications that aim to structurally estimate the parameters of a dynamic game. See, for example, Ericson and Pakes (1995).

each period are given by

$$u_i\left(a,(\theta,x)\right)=\begin{cases}\hat{u}_i(a,\theta)+x & \text{if } a_i=C \text{ and } \theta=L,\\ \hat{u}_i(a,\theta)-x & \text{if } a_i=D \text{ and } \theta=L,\\ \hat{u}_i(a,\theta)+2\cdot x & \text{if } a_i=C \text{ and } \theta=H,\\ \hat{u}_i(a,\theta)-2\cdot x & \text{if } a_i=D \text{ and } \theta=H,\end{cases}$$

where $\hat{u}_i(a,\theta)$ are the En-DPD stage-game payoffs in Table 2. The modification therefore adds an effective shock of $2\cdot x_t$ in the low state (or $4\cdot x_t$ in the high state) when contemplating a choice between C or D. The effect of the shock is static, as the draw next period x_{t+1} is independent, with an expected value of zero. The state-space swells from two payoff-relevant states in En-DPD to 22 here ($\{L,H\}\times X$, with the 11 states in X), where we will henceforth refer to this treatment as En-DPD-X.

Increasing the state-space leads to an increase in the set of admissible pure, symmetric Markov strategies. From four possibilities in the pivot, the increased state-space now allows for approximately 4.2 million Markov strategies. However, of the 4.2 million possibilities only one constitutes a symmetric MPE: cooperate at all states in $\{(L,x)\,|x\in X\}$, defect for all states in $\{(H,x)\,|x\in X\}$. The game therefore has the same effective MPE prediction as our pivot. Moreover, the efficient frontier of the extended-state–space game is (for the most part) unaltered, as are the set of simple SPEs.[29] Because of the strategic similarity to En-DPD, all the symmetric SPE that exist in the pivot have analogs here, while every efficient outcome is again supportable as an SPE using asymmetric history-dependent strategies. Importantly, given its focality in the pivot, joint cooperation can still be supported with a Markov trigger.

Examining the results for En-DPD-X in Figure 2, we see qualitatively similar average cooperation rates to those in the pivot. Comparing Figures 3(A) and (C), this similarity extends to the supergame level, though the slightly greater cooperation in both states for En-DPD-X is a little more apparent.[30] To make the comparison across treatments cleaner, the SFEM estimate uses the same strategy set as our previous treatments, and thus ignores strategies that condition on the shock $x_t\in X$.[31] The levels of equilibrium Markov play captured by the M_{CD} estimate are non-negligible, but compared to the less-complex pivot we actually see a decrease in its assessed weight. The largest difference between these two treatments is a substantial reduction of TfT in

[29]The sum of payoffs are maximized through any combination of (C,D)/(D,C) in the high state, unless $x_t\geq 3$, at which point (C,C) is superior.

[30]By the last five rounds, the average behavior depicted in Figure 2 for En-DPD-X is significantly more cooperative in both states, where this stems from greater stability in the response across the session.

[31]At the aggregate level, there is evidence of a correlation between the cooperation rate and the value of x in the high state. In the appendix, Figure 4 displays the cooperation rates for different values of x. Table 16 expands the SFEM analysis by including Markov and history-dependent strategies that condition on x. The main conclusions we present in the text are unaffected by this expansion.

favor of higher estimates for M_{CC}. This suggests that joint cooperation is more robust in En-DPD-X than the pivot, where some supergames are not triggering deviations after the first failure.

4.4. **Conclusions.** In the three treatments above we modify the pivot to examine whether some particular feature of our En-DPD parameterization are driving our results. In terms of first-order effects, all three robustness treatments continue to indicate a substantial selection of history-dependent behavior, if anything moving away from the MPE prediction and towards greater efficiency.

For our first two manipulations, the move towards efficiency comes despite *reducing the set of efficient outcomes.* The results suggest that the selection of history-dependent strategies over state-dependent ones is not solely driven by absolute-efficiency tradeoffs, but also the ease of coordination. In particular, in the two modifications, the changes do make it easier to coordinate on the MPE-trigger.

For our third manipulation, we examine a perturbation of the pivot with many more strategically relevant states. But the broad results are still very similar to those in the two-state pivot. This suggests a continuity in equilibrium selection with respect to the main strategic tensions of the dynamic game. The size of the state-space does not on its own increase the selection of MPE strategies. Though we perturb the game's presentation quite substantially, the outcomes in our En-DPD and En-DPD-X treatments are remarkably similar, reflecting the similar core game.

5. Changes to the Externalities

In the dynamic PD game treatments above there are two strategic considerations for each subject's chosen action. First, from a static point of view, their choice affects their partner's contemporaneous payoff. Second, from a dynamic perspective, their choice affects the transition across states, and hence their partner's future payoffs. Both strategic forces may lead subjects to cooperate more if they think inflicting these externalities on the other will affect future behavior. In this section we examine four new treatments that separate these two types of externality, to see how subjects' behavior responds to their absence. The first two treatments remove dynamic externalities, so that neither player's action choice affects future values for the state, holding constant the static externalities in the En-DPD pivot. The second treatment pair does the reverse: hold constant the pivot's dynamic externalities and remove the static externalities so neither player's choice affects the other's contemporaneous payoff.

5.1. **Removing Dynamic Strategic Externalities.**

Ex-DPD. To isolate the effects from dynamic externalities in En-DPD we change the transition rule. We fix the stage-games payoffs from the pivot (Table 2) so the static externalities are the same; however, we modify the state transition to remove any interdependence between the current state and the actions chosen last period. In this way we remove the dynamic externalities. For our first manipulation we choose an exogenous stochastic process for the new transition:

$$\psi(a,\theta)=\psi(\theta)=\begin{cases}3/5\cdot H\oplus 2/5\cdot L & \text{if } \theta=L\\ 4/5\cdot H\oplus 1/5\cdot L & \text{if } \theta=H.\end{cases}$$

The state evolves according to a Markov chain, which starts with certainty in the low state. If the state is low in any period, there is a 60 percent chance the game moves to high next period, and a 40 percent chance it remains in low. Given the present period is high, there is a 20 percent chance of a move to low next period, and an 80 percent chance it remains high.[32] Given this exogenous (Ex-) transition rule we label this dynamic PD treatment Ex-DPD.

All MPEs of a dynamic game with an exogenously evolving state are necessarily built-up from Nash profiles in the relevant stage games as the continuation value of the game is independent of the current actions. Because the stage-games in each state are PD games this leads to a unique MPE prediction: joint defection in both states. However, more-efficient SPE exist that allow for cooperation in the low state and (C,D)/(D,C) alternation in the high state.[33]

Looking at the experimental results for Ex-DPD, outcomes are starkly different from those where the state's evolution is endogenous. From Figure 2 it is clear that cooperation rates are much lower than the pivot, for both states. In the low state, the initial cooperation levels in the first period are 40–45 percent, where this falls across the supergame so that the overall low-state cooperation rate is closer to 30 percent. Cooperation in the high state is lower still, where average cooperation rates fall from 15 percent at the start of the session, to just under 10 percent in the final five supergames.

The reduced cooperation in Figure 2 is indicated at the supergame-level in Figure 3(E), where the large mass in the bottom-left corner is consistent with sustained defection in both states. This pattern is reflected too in the treatment's SFEM estimates in Table 4. The highest frequency is attributed to the MPE, M_{DD}, with an estimate of just under 60 percent. For those subjects who

[32]The Ex-DPD sessions were conducted after the En-DPD sessions were completed. The 60 percent and 80 percent probabilities were chosen to be close to aggregate state frequencies in the En-DPD sessions (66 and 83 percent, respectively).

[33]An asymmetric SPE that remembers whose turn it is to cooperate (defect) in high exists for $\delta=3/4$, given an M_{DD}-trigger on any deviation. History-dependent cooperation only in the low state can be sustained as a symmetric SPE with joint-defection in the high state at $\delta=3/4$, however, it is not an SPE to jointly cooperate in the high state, even with the worst-case M_{DD}-trigger on a deviation.

do attempt to support cooperation, the strategies used tend to be S_{DD}, reflecting a reversion to the MPE profile when cooperation is not successfully coordinated on.[34]

Removing the dynamic externalities dramatically shifts the observed behavior in the laboratory, leading to a collapse in cooperation. This finding suggests that the large cooperation rates that we document in the pivot treatment are largely related to subjects internalizing the consequences of their choices for the evolution of the state. We isolate this result further with our next treatment, which examines the extent to which the absence of *any* dynamics helps or hinders cooperation.

Ex-SPD. Our next modification goes further than Ex-DPD, so that there are no dynamics at all within a supergame. To do this we alter the transition rule to keep the game in the same fixed state for the entire supergame, so $\theta_{t+1} = \theta_t$ with certainty. Rather than a dynamic game, each supergame is now an infinitely repeated static PD (SPD) game, and we label this treatment Ex-SPD. To attain observations from subjects in both infinitely repeated stage games we make one additional change to the pivot, altering the starting state θ_1. For each supergame in Ex-SPD the starting period is the realization of the lottery, $3/5 \cdot H \oplus 2/5 \cdot L$. The chosen game therefore has the advantage of making the experimental environment and instructions similar to our other dynamic-game treatments (in terms of language, complexity and length).

Comparing aggregate-level results in Figure 2 it is clear that cooperation rates in Ex-SPD are higher for both states than for Ex-DPD. Because supergames are in a single fixed state, Figure 3(F) shows the results on separate axes. The figure indicates a large number of supergames with joint defection when the selected supergame state is high, but a larger degree of heterogeneity—and relatively more cooperation—when the supergame's state is low.[35] The comparison between Ex-SPD and Ex-DPD suggests that the possibility of the high state randomly occurring in Ex-DPD not only makes cooperation in the high state unlikely but also affects cooperation in the low state, where the incentives to cooperate are relatively higher.

In addition, we can compare behavior between the pivot and Ex-SPD. If we consider the Ex-SPD games in which $\theta = H$, the cooperation rate is slightly above 20 percent (last 5 cycles Figure 2). Meanwhile, for the same cycles, the cooperation rate when $\theta = H$ in En-DPD is higher than 50 percent. This indicates that the endogenous threat of returning to the low state upon deviation is crucial to supporting high levels of cooperation in En-DPD when $\theta = H$. In other words, it is not

[34]We also estimated strategy weights for this treatment adding the history-dependent strategy that supports cooperation only in the low state, described in footnote 33. The frequency estimate is 5.9 percent and is not significant. Subjects who aim to cooperate in this treatment try to cooperate in both states.

[35]There is almost no evidence of alternation between (C,D) and (D,C) outcomes in Ex-SPD when the state is high. This is consistent with previous findings in infinitely repeated PD and suggests that alternation documented for the pivot treatment is affected by dynamic incentives. In the pivot, a subject whose turn is to select C in high may oblige as the alternative choice of D would lead to the low state and possible permanently lower payoffs.

the case in En-DPD that cooperation would succeed at similar rates when $\theta = H$ if the threat of returning to the low state were absent.

SFEM estimates are presented by state in Table 4, and for this treatment we exclude from the estimation those strategies that condition differentially across states. When $\theta = H$, the frequency of always defect (here labeled M_{DD}) is comparable to the estimate for Ex-DPD. However, more-cooperative TfT strategies (both the standard and suspicious variety) are also selected, with aggregate frequencies close to 35 percent, higher than in Ex-DPD. The contrast with Ex-DPD behavior is starker in the low state. In this case, the frequency attributed to always defect (M_{DD}) is lower, where approximately two-thirds of the estimated strategies correspond to attempts to implement joint cooperation. The cooperation rates for both states in Ex-SPD are therefore in line with the larger experimental literature on infinitely repeated PD games, despite within-subject changes to the PD stage-game across the session.[36]

5.2. **Removing Static Strategic Externalities.** The previous subsection detailed what happens when we remove the pivot's dynamic externalities, but retain its static tensions. We now carry out the reverse exercise: turn off the static externalities, retaining the pivot's dynamic environment. Fixing the pivot's transition rule ψ—joint cooperation is required to transit from low to high, while anything but joint defection keeps the game in high—the next two treatments alter the stage-game payoffs, so that each player's static payoff is unaffected by the other's choice.[37] Removing static externalities means the stage-game is no longer a PD game, and we refer to this game instead as a dynamic common-pool (DCP) problem.[38] For greater comparability with the pivot, two separate parametrizations are used, with stage-games presented in Table 3.

[36]In infinitely repeated PDs, the basin of attraction of the grim-trigger (S_{DD}) helps predict cooperation. The basin of attraction of S_{DD} is the set of beliefs on the other's initial choice that would make S_{DD} optimal relative to M_{DD}. The low-state PD game has a basin of attraction for S_{DD} for any belief on the other also using S_{DD} above 0.24. In contrast, in the high-state game Grim is strictly dominated by playing always defect.

[37]The restriction is therefore that $u_i\left(\left(a_i, a_{-i}\right), \theta\right) = u_i\left(\left(a_i, a'_{-i}\right), \theta\right)$ for all $a_{-i}, a'_{-i} \in \mathcal{A}_{-i}$.

[38]In the DCP of Levhari and Mirman (1980) (LM from now on), there is a stock of fish that two fishermen can exploit. In each period, each agent independently decides how many fish to catch and their instantaneous payoffs are increasing in the number of fish they consume. In between periods the remaining fishes reproduce, determining the stock available in the next period. We highlight two features of this game. First, the transition is such that the lower the extraction in period t, the higher the remaining fish stock and, hence, the larger the available resource in the next period. Second, notice that (provided that no agent consumes more than fifty percent of the stock) the extraction of one agent in period t does not affect the other's agent payoff in period t. Our En-DCP treatments retain these features, but the main difference is in the number of states. In LM, the state is any non-negative number, which means that it can be feasible for low consumption levels to always lead to higher stocks in future periods. The efficient solution is symmetric and involves consumption levels lower than the symmetric MPE. With two states, as in our En-DCP treatments, the efficient solution cannot involve (C, C) and will instead be asymmetric in the high state. In order to keep the efficient frontier in the pivot and in the En-DCP treatments comparable we parametrize the pivot so that the efficient outcome is asymmetric in the high state.

TABLE 3. Dynamic Common Pool Treatments

(A) Markov Parametrization (En-DCP-M)

θ=Low

		2: C	2: D
1:	**C**	100,100	100, 125
	D	125,100	125,125

θ=High

		2: C	2: D
1:	**C**	130,130	130,190
	D	190,130	190 ,190

(B) Efficiency Parametrization (En-DCP-E)

θ=Low

		2: C	2: D
1:	**C**	100,100	100,125
	D	125,100	125,125

θ=High

		2: C	2: D
1:	**C**	130,130	130,280
	D	280, 130	280,280

Both parametrizations have the same payoffs in the low state: cooperation yields a payoff of 100, defection 125, regardless of what the other player chooses. The low-state payoff from selecting D corresponds to the pivot's temptation payoff, while the payoff from selecting C matches joint cooperation in the pivot. However, though selecting C in the low state involves a relative static loss of 25 it has a potential dynamic gain, the possibility of transiting to high next period if the other player also cooperates.

In the high state, we set the payoffs from choosing to cooperate at 130 in both parametrizations, which matches the high-state sucker's payoff in the pivot. The only difference between our two parametrizations is the payoff from choosing D in the high state. In the treatment we label "En-DCP-M" the payoff from defecting in high is set to 190, matching the pivot's joint-defection payoff. In the treatment we label "En-DCP-E" the payoff from defection is instead set to 280, matching the pivot's temptation payoff.

The En-DCP-M stage-game payoffs are chosen to match the payoffs attainable with the MPE (hence '-M') outcome in the pivot. The strategy M_{CD} in En-DCP-M yields exactly the same sequence of payoffs (and the same static/dynamic differences after any one-shot deviation) as the pivot. Although efficient outcomes still involve any combination of (C, D)/(D, C) in the high state, the payoffs realized from efficient paths here are lower than the pivot. To provide a control for this our En-DCP-E treatment's payoffs match the efficient (hence '-E') payoffs in the pivot. Conversely though, the payoffs from the most-efficient MPE are higher than in the pivot.

In both DCP treatments the most-efficient pure-strategy MPE uses M_{CD}, though M_{DD} also becomes a symmetric MPE. Efficient outcomes in both treatments are identical to the pivot and require asymmetric play.[39] If coordinated upon, taking turns cooperating and defecting in high can be supported as an SPE with a triggered reversion to either M_{CD} or M_{DD} in the En-DPD-M parametrization. Both A_{CD} and A_{DD} are SPE in En-DPD-M. However, this efficient outcome is only supportable as an SPE with an M_{DD} trigger in En-DPD-E (the strategy A_{DD}).

In terms of symmetry, the DCP treatments involve a change in the opposite direction from the efficiency manipulations presented in Section 4. Where those treatments lower the efficient frontier to make joint cooperation efficient, the DCP treatments fix the pivot's efficient outcomes and lower the value of symmetric cooperation. Joint cooperation is therefore *less* focal, and its static payoff is Pareto dominated by any action profile with defection. More so, joint-cooperation forever is not only less efficient than it was in the pivot, the symmetric MPE strategy M_{CD} is the Pareto-dominant symmetric SPE for our DCP treatments.

En-DCP-M treatment. The aggregate results in Figure 2 indicate reduced cooperation in both states relative to the pivot. However, the cooperation rate in the low state is still significantly greater than in the high state, particularly at the start of the supergame. At the history level Figure 3(G) shows a relatively large degree of variation across supergames, but with the largest mass concentrated at the bottom-right corner, consistent with the most-efficient MPE prediction, M_{CD}.

The SFEM estimates confirm the intuition from Figure 3(G), where the modal strategy is the most-efficient MPE with close to 30 percent of the mass. However, efficient asymmetric strategies that alternate in the high state do account for approximately a quarter of the data, suggesting a greater focus on them when the (slightly) less-efficient symmetric outcomes are removed. Just over 10 percent of the estimates reflect TfT, which as argued earlier can generate efficient asymmetric paths when it meets a complementary strategy. Relative to the pivot there is a large reduction in strategies implementing joint cooperation, where subjects avoid this pareto-dominated outcome.

En-DCP-E treatment. The patterns in our second common-pool parametrization have a starker match to the most-efficient MPE. The difference in average cooperation rates between the two states is larger than in En-DCP-M (Figure 2), where the largest mass of supergames are in the bottom-right corner of Figure 3(H). Looking at the SFEM results, the most-popular strategy by far is M_{CD}, with an estimated frequency close to two-thirds. History-dependent strategies that implement efficient outcomes are estimated at very low (and insignificant) frequencies. In fact, the only strategy showing a significant estimate involves reversion to M_{CD} when it (frequently) miscoordinates.

[39]Had the pivot game's efficient frontier involved *joint* cooperation we would not have been able to make a clear efficiency comparison with any DCP treatment. Instead, in our global experimental design, the efficient outcome in the pivot and En-DCP-E both require asymmetric high-state outcomes.

5.3. **Conclusion.** Behavior responds strongly to both dynamic and static externalities. Presenting data from four treatments that are similar to the pivot but removing one type of externality, we show that subjects' behavior responds with a greater selection of equilibrium Markov play than the pivot. The presence of both types of externality is therefore shown to be important for the selection of more-cooperative outcomes than the MPE.

Though our focus is on comparisons to the pivot, other interesting patterns emerge between treatments. Where we remove the dynamic externality from the pivot in Ex-DPD, conditional cooperation collapses, and the MPE of defection in both states becomes focal. However, when we remove the dynamics entirely, so that subjects face each stage game as a repeated game, then the cooperation rate in both states increases relative to Ex-DPD. Having an evolving state within the supergame makes it harder for subjects to cooperate. In particular, contrasting the findings in Ex-DPD and Ex-SPD suggests that the low cooperation rates in the high state of Ex-DPD may "contaminate" cooperation in the low state. Within a dynamic game the *interaction* between states is shown to be important, where diminished incentives to cooperate in future states can contaminate and reduce cooperation in current states, thus potentially pushing behavior towards the MPE.

For the dynamic common-pool treatments, a comparison of the two treatments suggests that even though both treatments are closer to the MPE, the evidence is much stronger in En-DCP-E. The reason for this a greater coordination on efficient asymmetric SPEs in the En-DCP-M treatments. However, as we increase the opportunity costs incurred from initiating efficient alternating cooperation in En-DCP-E—giving up 280 instead of 190 by cooperating first—this coordination on asymmetric outcomes disappears. It is possible that the absence of an efficient *symmetric* SPE is one of the drivers for the increased Markov play we observe, rather than the absence of static externalities. Though further research will likely separate between these forces with more precision, some evidence already exists. Vespa (2016) examines a dynamic common-pool game in which the state-space has no upper bound. This means that joint cooperation always leads to a higher payoff state and the efficient SPE is symmetric. While the incentives to cooperate are large in several treatments, modal behavior converges to a MPE and away from joint cooperation.

6. Discussion

6.1. **Summary of Main Results.** Our main experimental results indicate that:

Result 1 (History Dependence)**.** *Having a dynamic game does not necessarily lead to the selection of the MPE. Most subjects who do not use Markov strategies implement more efficient outcomes with history-dependent play.*

Evidence: Most behavior in En-DPD, En-DPD-CC, En-DPD-HT and En-DPD-X can be best described with more-efficient SPE strategies than any MPE profile. Though the symmetric MPE does very well at predicting some of our treatments (in particular Ex-DPD and En-DCP-E), the majority of our games are better explained via history-dependent strategies.[40]

Result 2 (Markov Selection)**.** *For subjects who use Markov profiles, the most-efficient MPE is the focal response.*

Evidence: In all treatments with endogenous transitions M_{CD} is the most-efficient MPE prediction. We find that this is the Markov strategy with the highest frequency in En-DPD, En-DCP-M, and En-DCP-E. In En-DPD-CC, En-DPD-HT and En-DPD-X the Markov strategy with the highest frequency is M_{CC}, but we note that this strategy is more-likely to be conflated with other history-dependent strategies.[41] In treatments with exogenous transitions, M_{DD} is the unique MPE and it is also the Markov strategy with the highest frequency.

Result 3 (Coordination and Efficiency)**.** *Reducing the static or dynamic temptations to deviate away from efficient symmetric SPE outcomes increase the selection of more-efficient cooperative outcomes.*

Evidence: In En-DPD-CC and En-DPD-HT we remove the possibility of supporting outcomes that compete with joint cooperation in terms with total payoffs (alternating between CD/DC in high). At the same time, in both treatments the highest possible equilibrium payoff is lower than in the pivot. In both cases, we find that cooperation increases, though more so for the dynamic modification in En-DPD-HT.

Result 4 (Perturbations)**.** *Adding exogenous states (shocks) does not lead to an increase in MPE play, and more cooperative outcomes are still common.*

Evidence: Our treatment with a richer state-space lead to differing rates of cooperation. Where we add exogenous non-persistent shocks to the payoffs each round (En-DPD→En-DPD-X) the aggregate observed behavior looks similar, if anything moving away from the MPE and towards higher-efficiency outcomes.

Result 5 (Response to Dynamic externalities)**.** *Behavior is sensitive to the removal of dynamic externalities, with a large reduction in cooperation rates and increase of MPE play.*

[40] In using the SFEM, our analysis focuses on allowing for errors in individual choices and estimating strategies that can rationalize the data. An alternative approach is to focus on an equilibrium concept that allows for errors. Goeree et al. (2016, chapter 5) detail a noisy equilibrium solution concept based on Markov strategies, a Markov Quantal Response Equilibrium (MQRE). In Appendix B we compute the MQRE for each of our treatments and contrast the prediction to the data. Our findings are in line with the results we report in the paper. The noisy MPE does relatively better in treatments where the SFEM reports a large proportion of strategies to be Markovian. Contrarily, where we find large proportions of history-dependent strategies, the MQRE can not accommodate the data. This suggests that subjects' choices in these treatments can not be rationalized as if they were playing an MPE with noise.

[41] See footnote 27 for a discussion of the identification problems between M_{CC} and more-forgiving strategies.

Evidence: Theory motivates that both static and dynamic externalities will drive whether an outcome is an equilibrium. In line with the theoretical prediction, there is a dramatic decrease in cooperation when dynamic externalities are not present (Ex-DPD) relative to the pivot. Moreover, we observe lower cooperation rates in either state of Ex-DPD than when each state is kept constant within the supergame (Ex-SPD), suggesting that the links between stage games in a dynamic game are important. Our results indicate that the selection of supra-MPE outcomes supportable by history-dependent play are affected by the presence of dynamic externalities (see also the increased cooperation in En-DPD-HT as we make this externality stronger).

Result 6 (Response to Static externalities)**.** *Behavior is sensitive to the removal of static externalities, with a reduction in cooperation rates and increase of MPE play.*

Evidence: Removing the static externality leads to a slight (En-DCP-M) or large (En-DCP-E) increase in the selection of equilibrium Markov strategies relative to the pivot. The presence of static externalities in En-DPD is shown to be an important component to the selection of history-dependent strategies in that treatment.

6.2. **Toward a Selection Index.** The above summarizes the main findings from our experiments. We now try to synthesize these results into a simple index for prediction that builds upon the more-extensive experimental literature on infinitely repeated games. That literature identifies two main determinants of history-dependent cooperative behavior (see the survey of Dal Bó and Fréchette 2014 for further details). First, whether or not cooperation can be supported as an SPE is predictive of outcomes. Second, and more fine-grained, the smaller the size of the basin of attraction (BA) for the always-defect strategy (M_D) relative to conditional cooperation (S_D, the grim trigger), the more likely cooperation is to emerge. The basin of attraction for M_D is a measure of the breadth of beliefs for the other player being a conditional cooperator that make M_D optimal relative to S_D. In other words, when a relatively low belief on the other cooperating is enough to make conditional cooperation attractive, then the basin for M_D is small, and cooperative behavior more likely to emerge. Notice that a situation in which a low belief on the other cooperating is enough to make conditional cooperation attractive is akin to a situation in which strategic uncertainty (not knowing what the other will do) is less crucial. That is, there is a connection between the basin of attraction of M_D and the notion of strategic uncertainty. As a rule of thumb, the literature offers the binary selection criterion: if the grim-trigger in the particular environment is risk-dominant (the M_D basin is smaller than a half) history-dependent cooperative SPE are predicted.

While our experiments were designed to investigate behavior across qualitative features of the game, a natural question given our results is whether predictive selection indices like the size of the BA can be generalized to dynamic environments. This leads to questions over which strategies are reasonable to construct a dynamic-game extension of the basin. For infinitely repeated PD games, the two strategies compared can be thought of as the MPE (M_D) and a symmetric strategy that supports the efficient outcome with a Markov trigger (S_D). But even in our simple dynamic games there are many potential MPEs and SPEs that might be used in the extension. Using the positive results in the last subsection we motivate the following: (i) The basin calculation should respond to *both* static and dynamic strategic externalities within the game, motivating extensions of the BA that integrate the entire dynamic game machinery into the calculation; (ii) Symmetric strategies are focal, and dynamic and static temptations measured again the best symmetric SPE are important for selection; (iii) Where the MPE is selected, the most-efficient MPE is the most-useful predictor; and (iv) Though we do find evidence for other strategies (for instance, tit-for-tat) trigger strategies that revert to the MPE on a deviation are common. Moreover, when we restrict the SFEM estimates to just these two strategies, the quality of fit reductions are not huge.[42]

The above motivates our focus on selection over two specific strategies: the dynamic-game Γ's most-efficient symmetric MPE (M_Γ); and the most-efficient symmetric outcome path sustainable as an SPE with reversion to M_Γ on any deviation (S_Γ). Why these two strategies? Three reasons: (i) our experimental results provide direct evidence that between them the MPE and MPE-trigger provide a large degree of rationalization for behavior in our games . (ii) Both strategies are theoretically tractable, and easy-to-calculate in applied settings. The MPE is the equilibrium object estimated in structural exercises, while the most-efficient outcome that can be supported by this MPE on a trigger is a relatively simple calculation (given the MPE and the one-deviation property). (iii) The choice of index dovetails with the previous repeated-games findings when the dynamic game is degenerate, where the theoretical idea underlying the calculation is familiar and easy to understand.

Our simple dynamic-game BA index is therefore $p^\star_\Gamma = p^\star(S_\Gamma, M_\Gamma; \Gamma)$: the probability of the other player choosing S_Γ that makes the agent indifferent between an ex ante choice of S_Γ or M_Γ. In our pivot game En-DPD, the two selected strategies would be S_{CD} and M_{CD}, and for $\delta = \frac{3}{4}$ the index calculation is $p^\star(S_{CD}, M_{CD}) = 0.246$, so that for all beliefs that the other will play S_{CD} above one-in-four, it is optimal to choose S_{CD} oneself.[43]

Given the index $p^\star_\Gamma$ we now examine how the index relates to empirical measures of MPE/supra-MPE behavior from our experimental sessions. In the infinitely-repeated game literature the focal outcome measure is the first-period cooperation rate in supergames. However, this measure does

[42]See Table 16 in the appendix for details and estimates.

[43]The calculation leads to the following normal-form representation for the row player's discounted-average payoff (where $\pi^H_{M_{CD}} = \frac{1}{1+\delta} \cdot 190 + \frac{\delta}{1+\delta} \cdot 100$ is the high-state payoff under the MPE M_{CD}):

	S_{CD}	M_{CD}
S_{CD}	$(1-\delta) \cdot 100 + \delta \cdot 200$	$(1-\delta) \cdot 100 + (1-\delta)\delta \cdot 130 + \delta^2 \cdot \pi^H_{M_{CD}}$
M_{CD}	$(1-\delta) \cdot 100 + (1-\delta)\delta \cdot 280 + \delta^2 \cdot \pi^H_{M_{CD}}$	$(1-\delta) \cdot 100 + \delta \cdot \pi^H_{M_{CD}}$

Note that the first-period action payoff is the same regardless of the cell and so will not affect the basin calculation.

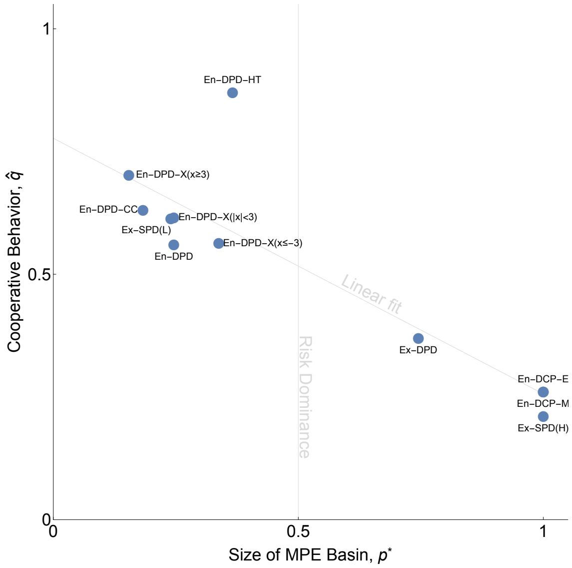

(A) Cooperation at divergence

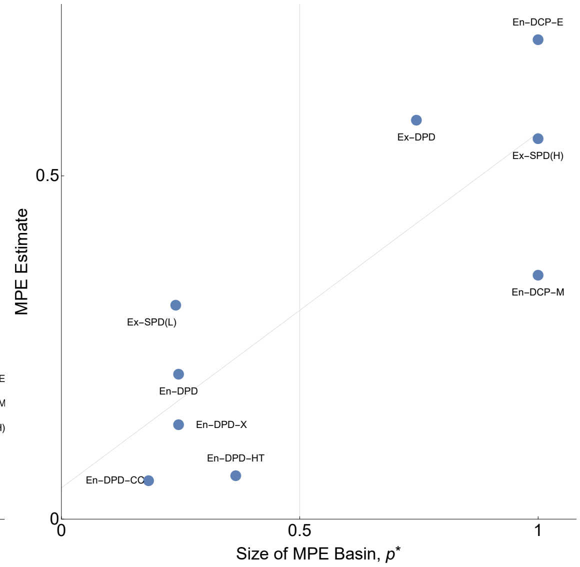

(B) SFEM Estimate

FIGURE 1. Basin of Attraction for the MPE

Note: The figures horizontal axis shows the size of the Basin of Attraction for the MPE relative to the most-efficient symmetric SPE with an MPE trigger, $p^\star\left(S_\Gamma, M_\Gamma; \Gamma\right)$. The vertical axis in panel (A) presents the average cooperation rate in the period where the two strategies in the basin calculation diverge $\hat{q}$ (period two for all treatments except Ex-DPD and Ex-SPD with low and high starting states, where coordination is resolved after the first period). The vertical axis in panel (B) presents the SFEM estimates for the MPE from Table 4 .

not capture supra-MPE behavior in many of our dynamic-game settings. For example, in the pivot treatment, both the MPE and SPE strategies are predicted to cooperate in the first (low-state) period. First-round cooperation will therefore not be informative on the selection margin we are interested in. Instead we focus on two distinct measures: (i) the raw cooperation rate $\hat{q}_\Gamma$ in the first instance where M_Γ and S_Γ are predicted to diverge; and (ii) the estimated weight $\phi(M_\Gamma)$ on the MPE strategy from our SFEM mixture model.

For the pivot, the strategies S_{CD} and M_{CD} both cooperate in round one and then choose differing actions in the high state in the second-round, C and D, respectively. Using the data from the last five supergames in the pivot sessions, 56.2 percent of our subject-supergames have paths consistent with $\left(\left(C, C\right), L\right), \left(C, \cdot, H\right)$ and so for the pivot the basin-behavior pair $(p^\star, \hat{q})$ is given by $(0.246, 0.562)$.[44] Examining Table 4, the basin-MPE estimate pair $\left(p^\star, \phi\left(M_\Gamma\right)\right)$ is given by $(0.246, 0.212)$. Figures 1(A) and 1(B) provide plots for each $(p^\star, \hat{q})$ and $\left(p^\star, \phi\left(M_\Gamma\right)\right)$ pair across

[44]In this way, we are more conservative in ascribing behavior as cooperative. A player may have been cooperative in period one, but was defected on, and will be counted as a zero for our supra-MPE measure. The outcome $\hat{q}_\Gamma$ only positively incorporates behavior with a *perfectly* consistent path with both S_Γ and M_Γ up to the predicted point of divergence, and that then takes the prescribed S_Γ action.

our treatment set (and some sub-treatments where the prediction differs) illustrating the basin's predictive power.[45]

As illustrated by the two panels in Figure 1, the theoretical basin calculation organizes both the MPE estimates and the supra-MPE cooperation well. For example, given the 12 data points represented in the first panel, an OLS regression indicates that cooperation rate at divergence is significantly related to MPE basin size (99 percent confidence). Moreover, the figure also shows that the easier-to-interpret binary criterion for the MPE assumption is also predictive, whether or not the MPE risk dominates the most-efficient Markov-trigger SPE.

Given the analytically tractable form of the index, the relationship we find should make us more positive on discovering more fundamental regularities to selection. This is particularly so, given the qualitatively very different treatments. With a proposal for an index, future experimental research can design an experiment to explicitly test it. However, caution with respect to the domain and details of the index are obviously warranted. For example, the index is rather crude relative to some of our findings. We document that treatments in which an efficient outcome is more focal (for example, En-DPD-HT) it easier for subjects to coordinate on efficient outcomes, but the index does not capture how *focal* efficient outcomes are. Thus, while our index is relatively simple, there may be meaningful ways to expand upon it that can account for further nuances of the games.

7. Conclusion

Our paper explores a set of eight dynamic games under an infinite-time horizon. While many applications of dynamic games focus on Markov-perfect equilibria, our results suggest more-nuanced selection. In our results the selection of state-dependent strategies responds to strategic features of the game.

[45] While in some treatments the same basin calculation between S_{CD} and M_{CD} is used (in particular En-DPD-CC and En-DPD-HT) in others the basin calculation has to change. For instance, though the strategies over which selection is calculated stay the same in En-DPD-X the riskiness of coordination on S_{CD} is influenced by the second-period shock x_2. For negative values of x_2, the index is higher indicating that cooperation is less likely, while the opposite happens for positive values of x_2. In Figure 1 we aggregate the shocks into three categories ($x_2 \leq -3$, $-3 < x_2 < 3$, and $3 \leq x_2$), and plots the basin-behavior pairs. For Ex-DPD the basin calculation shifts to account for changes in both the MPE and most-efficient symmetric SPE. The MPE prediction in this game shifts to M_{DD} because of the exogenous transitions. Additionally, S_{DD} is not an SPE, where the basin calculation calls for the best symmetric SPE. Instead we use the symmetric trigger that supports cooperation in low and defection in high with a reversion to M_{DD} on any deviation. In the Ex-SPD treatment where the low state is initially selected the basin calculation is the standard infinitely-repeated PD game calculation $p^\star\left(S_D, M_D\right)$. Finally, in three treatments the best-*symmetric* SPE *is* the most-efficient MPE, in which case the basin for the MPE is the full set of beliefs, with $p^\star = 1$. This is true for our two DCP treatments (which have identical rates of high-state cooperation in round two, so the plotted data-points are coincident) and the Ex-SPD supergame that starts in the high state. Additionally, for the Ex-DPD and Ex-SPD games, coordination issues are resolved in round one; Figure 1 reflects this as $\hat{q}_\Gamma$ is the period-one cooperation rate for these treatments.

Our core treatments are simple two-state extensions of the infinitely repeated prisoner's dilemma, and we find that most behavior is consistent with history-dependent strategies that aim to achieve greater efficiency than the MPE prediction. The results from our dynamic PD treatments illustrate a richness in subject behavior, where both the state and recent history affect choice. However, where we weaken the strategic externalities (in both static and dynamic senses) behavior in the game becomes more state-dependent.

Extending our findings, future research can help delineate the boundaries between dynamic games where the MPE is or is not likely to be selected. Many first-order questions remain open. For instance, in dynamic game environments little is known about how equilibrium selection responds to the importance of the future (via the discount factor). Similarly, greater experimentation over the size of the action space, or the number of other players may help us understand the role of strategic uncertainty in equilibrium selection. The laboratory offers a powerful tool to induce and control strategic environments, and measure human behavior within them. This tool can be particularly useful for dynamic games: environments where equilibrium theory will generically result in a multiplicity of prediction, but where applications require specificity and tractability. Experiments can help not only in validating the particular settings where MPE assumptions represent behavior; but also for those settings where it seems unlikely they can offer data-based alternatives.

References

Acemoglu, Daron and James A Robinson, "A theory of political transitions," *American Economic Review*, 2001, pp. 938–963.

Aghion, Philippe, Christopher Harris, Peter Howitt, and John Vickers, "Competition, imitation and growth with step-by-step innovation," *The Review of Economic Studies*, 2001, *68* (3), 467–492.

Bajari, Patrick, C Lanier Benkard, and Jonathan Levin, "Estimating dynamic models of imperfect competition," *Econometrica*, 2007, *75* (5), 1331–1370.

Battaglini, M. and S. Coate, "Inefficiency in Legislative Policymaking: A Dynamic Analysis," *The American Economic Review*, 2007, pp. 118–149.

______, S. Nunnari, and T. Palfrey, "The Dynamic Free Rider Problem: A Laboratory Study," *mimeo*, 2014.

Battaglini, Marco, Salvatore Nunnari, and Thomas R Palfrey, "Legislative bargaining and the dynamics of public investment," *American Political Science Review*, 2012, *106* (02), 407–429.

Benchekroun, Hassan, Jim Engle-Warnick, and Dina Tasneem, "An Experimental Study of a Common Property Renewable Resource Game in Continuous Time," *CIRANO-Scientific Publications 2014s-09*, 2014.

Bergemann, D. and J. Valimaki, "Dynamic common agency," *Journal of Economic Theory*, 2003, *111* (1), 23–48.

Bó, Pedro Dal and Guillaume R Fréchette, "The evolution of cooperation in infinitely repeated games: Experimental evidence," *The American Economic Review*, 2011, *101* (1), 411–429.

Coles, M.G. and D.T. Mortensen, "Dynamic Monopsonistic Competition and Labor Market Equilibrium," *mimeo*, 2011.

Dal Bó, Pedro and Guillaume R Fréchette, "On the Determinants of Cooperation in Infinitely Repeated Games: A Survey," 2014.

Dixit, A., G.M. Grossman, and F. Gul, "The dynamics of political compromise," *Journal of Political Economy*, 2000, pp. 531–568.

Dutta, P.K. and R. Radner, "Population growth and technological change in a global warming model," *Economic Theory*, 2006, *29* (2), 251–270.

Ericson, R. and A. Pakes, “Markov-perfect industry dynamics: A framework for empirical work,” *The Review of Economic Studies*, 1995, *62* (1), 53.

Fréchette, Guillaume R and Sevgi Yuksel, “Infinitely Repeated Games in the Laboratory: Four Perspectives on Discounting and Random Termination,” February 2013. NYU working paper.

Fudenberg, D., D.G. Rand, and A. Dreber, “Slow to anger and fast to forgive: Cooperation in an uncertain world,” *American Economic Review*, 2010.

Goeree, Jacob A., Charles A. Holt, and Thomas R. Palfrey, *Quantal Response Equilibrium: A Stochastic Theory of Games* 2016.

Hörner, J. and L. Samuelson, “Incentives for Experimenting Agents,” *mimeo*, 2009.

Kloosterman, A., “An Experimental Study of Public Information in Markov Games,” *mimeo*, 2015.

Laibson, D., “Golden Eggs and Hyperbolic Discounting,” *Quarterly Journal of Economics*, 1997, *112* (2), 443–477.

Levhari, David and Leonard J Mirman, “The great fish war: an example using a dynamic Cournot-Nash solution,” *The Bell Journal of Economics*, 1980, pp. 322–334.

Mailath, George J and Larry Samuelson, *Repeated games and reputations*, Vol. 2, Oxford university press Oxford, 2006.

Maskin, Eric and Jean Tirole, “A theory of dynamic oligopoly, I: Overview and quantity competition with large fixed costs,” *Econometrica: Journal of the Econometric Society*, 1988, pp. 549–569.

______ **and** ______ , “Markov perfect equilibrium: I. Observable actions,” *Journal of Economic Theory*, 2001, *100* (2), 191–219.

Rojas, Christian, “The role of demand information and monitoring in tacit collusion,” *RAND Journal of Economics*, 2012, *43* (1), 78–109.

Rubinstein, A. and A. Wolinsky, “Decentralized trading, strategic behaviour and the Walrasian outcome,” *The Review of Economic Studies*, 1990, *57* (1), 63.

Saijo, T., K. Sherstyuk, N. Tarui, and M. Ravago, “Games with Dynamic Externalities and Climate Change Experiments,” *mimeo*, 2014.

Salz, Tobias and Emanuel Vespa, “Estimating Dynamic Games of Oligopolistic Competition: An Evaluation in the Laboratory,” 2016. UCSB working paper.

Sherstyuk, Katerina, Nori Tarui, and Tatsuyoshi Saijo, “Payment schemes in infinite-horizon experimental games,” *Experimental Economics*, 2013, *16* (1), 125–153.

Vespa, Emanuel, “An Experimental Investigation of Strategies in Dynamic Games,” 2016. UCSB working paper.

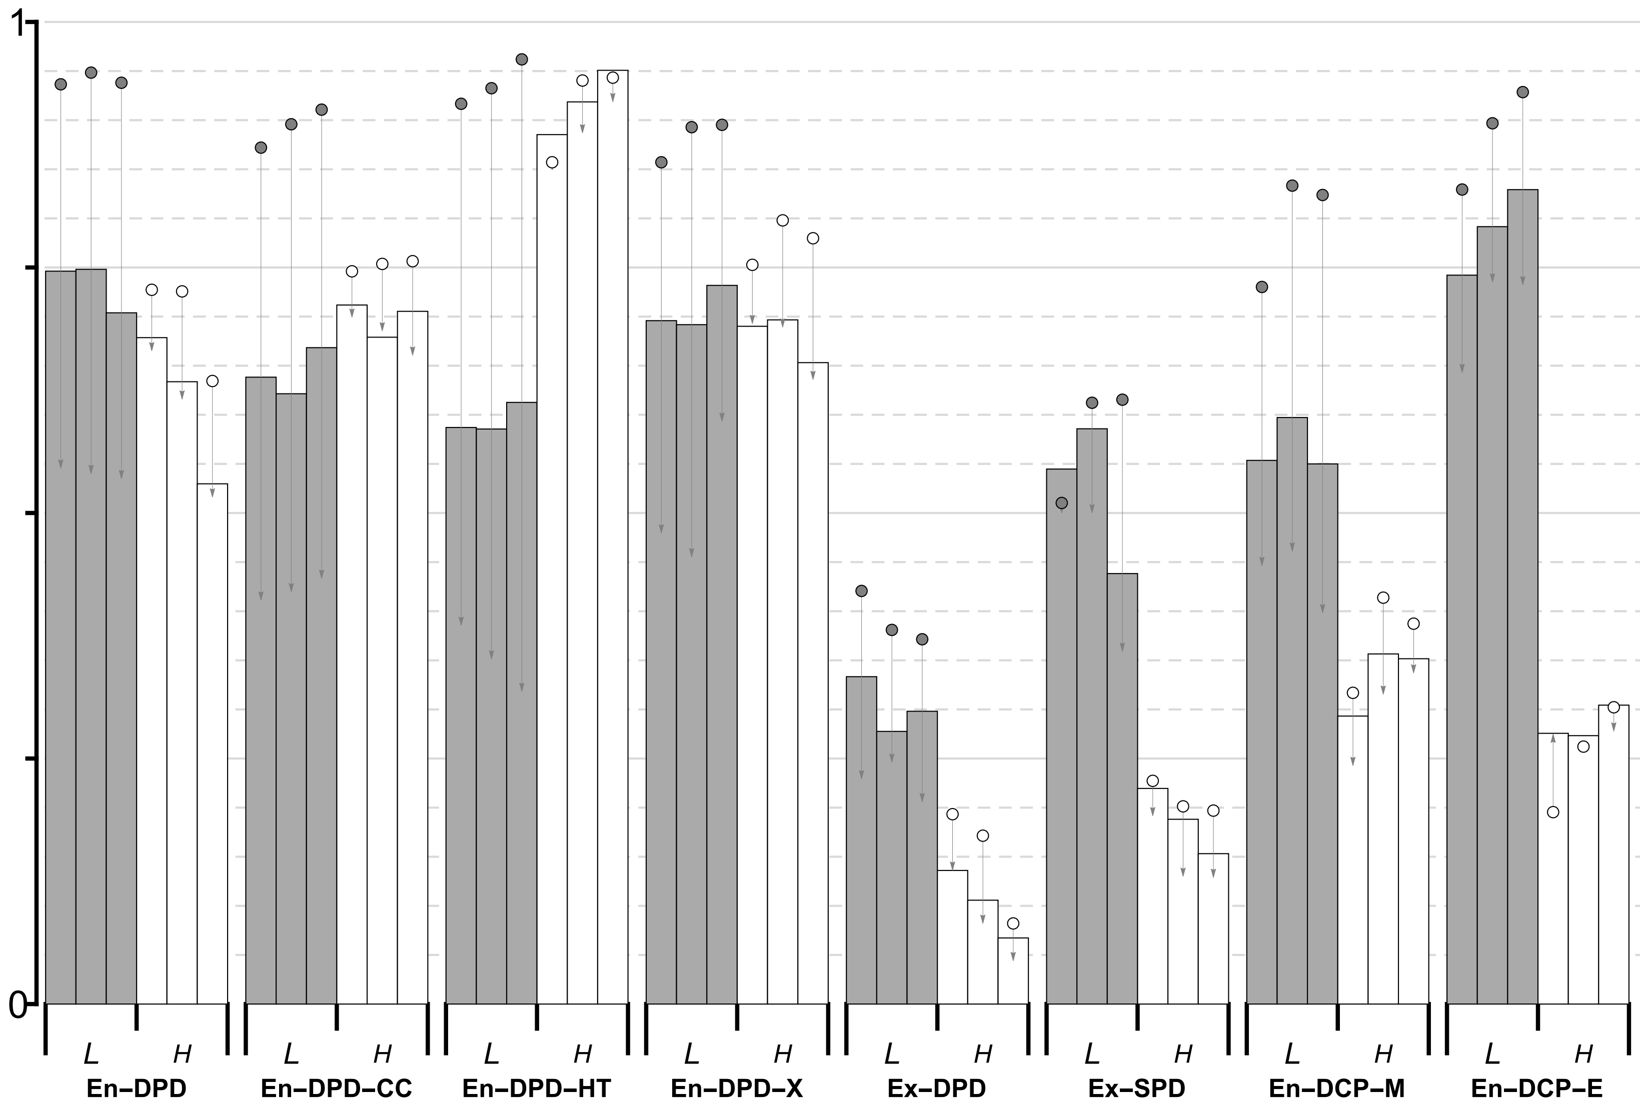

FIGURE 2. Cooperation Rates, by treatment/state/supergame-block

Note: The unit of observation is a period. Cooperation rates are given in blocks of five supergames, where the first bar in each sequence illustrates cooperation rates in supergames 1–5, the second supergames 6–10, and the last supergames 11–15. Circular points indicate the cooperation rate in period one of the supergame for Low states (all supergames), and period two for the high states (only those supergames which enter the high state in period two), except for Ex-SPD, where both circles shows show period one cooperation. Arrows point to the final cooperation rate (last two periods in a supergame) in each state.

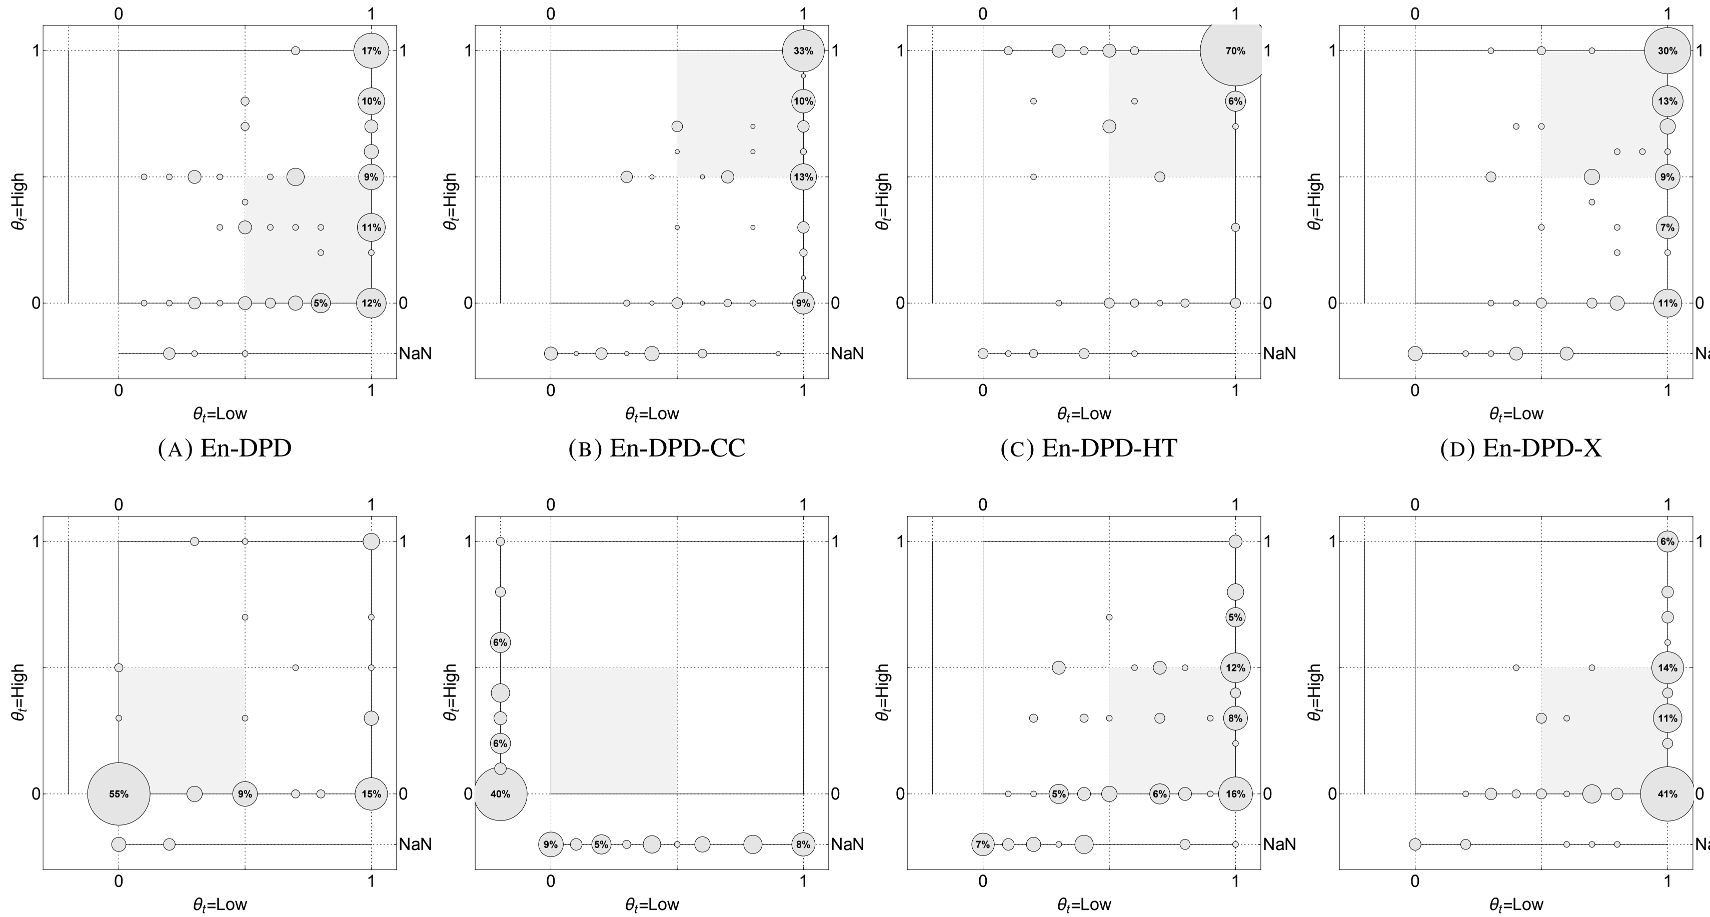

(A) En-DPD (B) En-DPD-CC (C) En-DPD-HT (D) En-DPD-X

(E) Ex-DPD (F) Ex-SPD (G) En-DCP-M (H) En-DCP-E

FIGURE 3. Histories (last five supergames)

Note: The unit of observation is a history: the choices of a subject in a supergame. The data are represented in each Figure on an 11 by 11 grid, so that for example, a cooperation rate of 97 percent in one state is represented as 100 percent.

TABLE 4. Strategy Frequency Estimation Method Output: Last Five Supergames

Strategies	En-DPD	En-DPD-CC	En-DPD-HT	En-DPD-X	Ex-DPD	Ex-SPD ($\theta = L$)	Ex-SPD ($\theta = H$)	En-DCP-M	En-DCP-E
Markov									
M_{CC}	0.117	0.173**	0.251**	0.347***	0.000	0.097*	0.023	0.068	0.092
	(0.072)	(0.076)	(0.116)	(0.091)	(0.013)	(0.055)	(0.030)	(0.053)	(0.059)
M_{DD}	0.024	0.039	0.023	0.027	0.582***	0.312***	0.555***	0.077	0.048
	(0.030)	(0.039)	(0.017)	(0.034)	(0.145)	(0.115)	(0.101)	(0.061)	(0.038)
M_{CD}	0.212*	0.057	0.041	0.138*	0.073*			0.279*	0.651***
	(0.127)	(0.063)	(0.033)	(0.081)	(0.041)			(0.158)	(0.181)
M_{DC}	0.000	0.000	0.000	0.000	0.000			0.063	0.000
	(0.003)	(0.010)	(0.026)	(0.015)	(0.000)			(0.043)	(0.002)
History-dependent									
S_{DD}	0.106	0.227*	0.479***	0.069	0.180*	0.255*	0.074	0.039	0.000
	(0.095)	(0.119)	(0.133)	(0.066)	(0.109)	(0.150)	(0.064)	(0.048)	(0.009)
S_{CD}	0.206**	0.075	0.000	0.245**	0.045			0.070	0.088*
	(0.085)	(0.067)	(0.043)	(0.110)	(0.033)			(0.057)	(0.052)
TfT	0.254***	0.304***	0.182	0.089	0.032	0.280***	0.121	0.139*	0.069
	(0.082)	(0.093)	(0.125)	(0.063)	(0.052)	(0.094)	(0.076)	(0.080)	(0.078)
$STfT$	0.023	0.016	0.000	0.021	0.065	0.055	0.227	0.000	0.000
	(0.027)	(0.020)	(0.002)	(0.031)	(0.062)			(0.008)	(0.002)
A_{DD}	0.059	0.046	0.024	0.029	0.022			0.164	0.051
	(0.039)	(0.041)	(0.029)	(0.034)	(0.043)			(0.111)	(0.050)
A_{CD}	0.000	0.065	0.000	0.035	0.000			0.102	0.000
β	0.826	0.869	0.940	0.828	0.947	0.893	0.921	0.805	0.870
Wald Tests on Markov Strategies									
Mass consistent with any MPE	0.212	0.057	0.064	0.138	0.582	0.312	0.555	0.358	0.699
Null hypothesis:									
Markov Strategies=0	0.315	0.150	0.025	0.000	0.000	0.004	0.000	0.011	0.000
Markov Strategies=1	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
MPE=0	0.095	0.367	0.146	0.089	0.000	0.007	0.000	0.006	0.000
MPE=1	0.000	0.000	0.000	0.000	0.004	0.000	0.000	0.000	0.000
# Observations	1,260	1,568	1,358	1,204	1,050	454	764	1,232	1,218

Note: Bootstrapped standard errors in parentheses. Level of Significance: ***-1 percent; **-5 percent; *-10 percent. Wald Tests on Markov Strategies: The row 'Mass consistent with any MPE' reports the sum of the estimated mass corresponding to a Markov strategy that is sub-game perfect. (See Table 7 for details on which Markov strategy is an SPE by treatment.) The remaining rows report the p-value corresponding to the null hypotheses that the sum of the coefficients of all pure Markov strategies equals 0 (Markov Strategies=0), equals 1 (Markov Strategies=1), and the null hypotheses that the sum of the coefficient of Markov strategies that are sub-game perfect equals 0 (MPE=0) and equals 1 (MPE=1).

APPENDIX A. FOR ONLINE PUBLICATION: SUPPLEMENTARY MATERIAL: FIGURES AND TABLES

Tables 5 and 6 present the stage games for the En-DPD-CC and En-DPD-X treatments, respectively.

TABLE 5. En-DPD-CC Stage Games

θ=Low

		2: C	2: D
1:	**C**	100,100	30, 125
	D	125, 30	60,60

θ=High

		2: C	2: D
1:	**C**	200, 200	130, 250
	D	250, 130	190, 190

TABLE 6. En-DPD-X Stage games

$\theta = (\textbf{Low}, x)$

		2: C	2: D
1:	**C**	100+x,100+x	30-x, 125+x
	D	125-x, 30+x	60-x,60-x

$\theta = (\textbf{High}, x)$

		2: C	2: D
	C	200+2x, 200+2x	130+2x, 280-2x
	D	280-2x, 130+2x	190-2x, 190-2x

TABLE 7. Summary of theory predictions - Is each strategy an SPE of the game?

Strategies	**En-DPD**	**En-DPD-CC**	**En-DPD-HT**	**En-DPD-X**	**Ex-DPD**	**Ex-SPD** ($\theta = L$)	**Ex-SPD** ($\theta = H$)	**En-DCP-M**	**En-DCP-E**
Markov									
M_{CC}	No	No	No	No	No	No	No	No	No
M_{DD}	No	No	Yes	No	Yes	Yes	Yes	Yes	Yes
M_{CD}	Yes	Yes	Yes	Yes	No	-	-	Yes	Yes
M_{DC}	No	No	Yes	No	No	-	-	No	No
History-dependent									
S_{DD}	No	No	Yes	No	No	Yes	No	No	No
S_{CD}	Yes	Yes	Yes	Yes	No	-	-	No	No
TfT	No	No	No	No	No	No	No	No	No
$STfT$	No	No	No	No	No	No	No	No	No
A_{DD}	No	No	No	No	No	-	-	Yes	Yes
A_{CD}	Yes	Yes	No	Yes	No	-	-	Yes	No

TABLE 8. Cooperation rates by state (Last 5 supergames)

Treatment	θ_t =**Low**			θ_t =**High**		
	Mean	**(std. err)**		**Mean**	**(std. err)**	
En-DPD	0.796	(0.035)	–	0.489	(0.045)	–
En-DPD-CC	0.794	(0.036)		0.674	(0.042)	$(\star\star\star)$
En-DPD-HT	0.832	(0.050)		0.979	(0.010)	$(\star\star\star)$
En-DPD-X	0.856	(0.036)	$(\star)$	0.635	(0.055)	$(\star\star\star)$
Ex-DPD	0.189	(0.059)	$(\star\star\star)$	0.012	(0.007)	$(\star\star\star)$
Ex-SPD†	0.406	(0.062)	$(\star\star\star)$	0.079	(0.024)	$(\star\star\star)$
En-DCP-M	0.638	(0.047)	$(\star\star\star)$	0.245	(0.041)	$(\star\star\star)$
En-DCP-E	0.946	(0.021)	$(\star\star)$	0.187	(0.047)	$(\star\star\star)$

Note: Unit of observation: choice of a subject in a period. Figures reflect predicted cooperation rates for the median subject (subject random-effect at zero) attained via a random-effects probit estimate over the last five cycles with just the state as a regressor. Statistical significance is given for differences with the pivot En-DPD, except for: †- Statistical significance here given relative to Ex-DPD

Further analysis at the aggregate level. Table 8 presents tests on whether the cooperation rates by state and treatment in Figure 2 are statistically different from the pivot. The predicted cooperation rates are obtained after estimating a random-effects probit with a dummy variable for cooperation in the left-hand-side, and a constant and a state dummy on the right-hand side.

Table 9 performs a robustness check on the estimates of Table 8. The table reports the estimates of a linear probability model with the same dependent variable, but an additional set of controls and standard errors that are clustered at the session level. Each treatment presents estimates relative to the pivot, so that the Treatment dummy takes value 1 if the observation corresponds to that treatment and 0 if it belongs to the pivot. There is also a state dummy and the interaction between state and treatment dummies. Finally, there is a set of dummy variables for the included supergames.

Table 11 provides marginal effect estimates from a probit examining the decision to cooperate after the subject has already chosen supra-MPE cooperation. For the majority of our treatments we look at those subjects (i, matched to j) at the history $h_3(a_{j2}) = ((C, C, L), (C, a_{j2}, H))$, where their partner has chosen $a_{j2} \in \{C, D\}$. The probit assesses the extent to which the other player's action choice affects i's decision to continue cooperating in the high state, where we additionally control for the supergame (1 to 15) within the session.

Tables 12 and 13 report the most frequently observed evolution of the state and sequences of actions, respectively.

TABLE 9. Cooperation relative to Pivot Treatment (Last 5 Supergames): Panel Regression

	En-DPD-CC	En-DPD-HT	En-DPD-X	Ex-DPD	Ex-SPD	En-DCP-M	En-DCP-E
Constant	0.956***	0.956***	0.963***	0.987***	0.971***	0.984***	0.960***
	(0.028)	(0.027)	(0.025)	(0.035)	(0.026)	(0.023)	(0.026)
Treatment	-0.027	0.023	-0.042	-0.566***	-0.331***	-0.114***	-0.009
	(0.048)	(0.036)	(0.050)	(0.112)	(0.106)	(0.032)	(0.040)
State	-0.331***	-0.333***	-0.330***	-0.335***	-0.336***	-0.329***	-0.332***
	(0.057)	(0.057)	(0.058)	(0.057)	(0.056)	(0.059)	(0.057)
State $\times$ Treatment	0.136*	0.294***	0.163***	0.067	-0.052	-0.169***	-0.318***
	(0.075)	(0.067)	(0.059)	(0.089)	(0.061)	(0.065)	(0.089)
Supergame 12	-0.017	-0.011	-0.010	-0.015	-0.017	-0.027	-0.031**
	(0.031)	(0.023)	(0.022)	(0.038)	(0.031)	(0.026)	(0.015)
Supergame 13	-0.024*	-0.029*	-0.035	-0.039	-0.069***	-0.074***	-0.011
	(0.013)	(0.015)	(0.025)	(0.033)	(0.026)	(0.027)	(0.023)
Supergame 14	-0.017	-0.026	-0.032	-0.087***	-0.042	-0.076***	-0.042*
	(0.039)	(0.029)	(0.036)	(0.032)	(0.028)	(0.020)	(0.022)
Supergame 15	-0.034	-0.024	-0.047***	-0.104*	-0.037	-0.055***	-0.026
	(0.021)	(0.021)	(0.016)	(0.062)	(0.032)	(0.017)	(0.025)
# Obs.	910	800	774	740	606	748	788

Note: Unit of observation: choice of a subject in a selected period. Selected periods include: first period in the low state and first period in the high state. Treatment is a dummy variable that takes value 1 for the treatment corresponding to the column and zero for the pivot (En-DPD) treatment. State is a dummy variable that takes value 1 if the state is high, 0 if the state is low. State $\times$ Treatment is the interaction of the state and treatment dummies. Each 'Supergame' variable is a dummy variable that takes value 1 for the corresponding supergame, 0 otherwise. The dependent variable takes value 1 if the subject cooperated, 0 if the subject defected. The data includes all period 1 choices and for all treatments but En-SPD all period 2 choices when the state for that period is High. Each column reports the results of a random effects linear probability model and standard errors (reported between parentheses) are clustered at the session level. Level of Significance: ***-1 percent; **-5 percent; *-10 percent.

TABLE 10. Differences between initial and subsequent period Cooperation Rates

Treatment	$\theta =$**Low**			$\theta =$**High**		
	$\Delta\Pr\{C\}$	**(Std. Err)**		$\Delta\Pr\{C\}$	**(Std. Err)**	
En-DPD	0.498	(0.075)	$(\star\star\star)$	0.213	(0.046)	$(\star\star\star)$
En-DPD-CC	0.520	(0.066)	$(\star\star\star)$	0.135	(0.044)	$(\star\star\star)$
En-DPD-HT	0.867	(0.090)	$(\star\star\star)$	0.006	(0.014)	
En-DPD-X	0.256	(0.065)	$(\star\star\star)$	0.256	(0.055)	$(\star\star\star)$
Ex-DPD	0.124	(0.050)	$(\star\star\star)$	0.049	(0.022)	$(\star\star)$
Ex-SPD	0.286	(0.068)	$(\star\star\star)$	0.040†	(0.021)	$(\star)$
En-DCP-M	0.421	(0.053)	$(\star\star\star)$	0.069	(0.039)	$(\star)$
En-DCP-E	0.084	(0.039)	$(\star\star)$	0.040	(0.034)	

Note: Unit of observation: choice of a subject in a period. Figures reflect predicted marginal effect $\Delta\Pr\{C\}$=$\Pr\{C|$ Initial Period$,\theta\}-\Pr\{C|$Subsequent Period$,\theta\}$ for the initial play dummies for the median subject (subject random effect at zero) attained via a random-effects probit estimate over the last five cycles (regressors are state dummies and dummies for Low & Period One and High & Period 2;. Statistical significance is relative to zero. †-For Ex-DPD we define the initial level with a High & Period 1 dummy.

TABLE 11. History Dependence Probits

	En-DPD	En-DPD- CC	En-CP-L	En-CP-H	En-DPD-X	Ex-DPD	Ex-SPD	Ex-SPD
History Selection	$h_3 = ((C,C,L),(C,a_{j2},H))$					$h_2 = (C,a_{j2},L)$		$h_2 = (C,a_j,H)$
$\mathbf{1}\{a_{j2}=C\}$	0.425***	0.499***	0.198	0.150	0.349***	0.622***	0.424***	0.499***
	(0.076)	(0.0712)	(0.124)	(0.094)	(0.079)		(0.111)	(0.124)
Supergame	-0.019***	0.001	-0.006	-0.009	-0.006		-0.020***	0.018**
	(0.005)	(0.003)	(0.011)	(0.008)	(0.004)		(0.007)	(0.008)
Base Prob.	0.424	0.437	0.321	0.313	0.521		0.519	0.371
	(0.070)	(0.067)	(0.073)	(0.124)	(0.0743)		(0.106)	(0.086)
Supra-MPE Coop	0.700	0.751	0.376	0.254	0.769	0.389	0.584	0.242
Supra-MPE Coop$\times\beta_{\mathbf{1}\{a_{j2}=C\}}$	0.298	0.375	0.075	0.038	0.268	0.242	0.248	0.121
N (*S*)	389 (38)	506 (49)	151 (30)	127 (24)	377		146 (36)	92 (24)

Note: The *Supra-MPE Coop* row reports the cooperation round-state in the first round where the MPE does not prescrive cooperation. We interact this with the estimate for history dependence to provide an estimate for the fraction of that cooperation that switches after a single deviation.

TABLE 12. Path for the State: Last Five Supergames

Sequence	En-DPD	En-DPD-CC	En-DPD-HT	En-DPD-X	Ex-DPD	Ex-SPD	En-DCP-M	En-DCP-E
LLLLL	0.038	0.121	0.057	0.086	0.048	0.371	0.200	0.048
	(0.004)	(0.009)	(0.005)	(0.008)	(0.004)	(0.023)	(0.016)	(0.004)
LHLHL	0.038	0.014	0.029	0.019	0.00		0.114	0.362
	(0.004)	(0.001)	(0.003)	(0.002)	–		(0.010)	(0.023)
LHHHH	0.476	0.607	0.733	0.600	0.371		0.19	0.267
	(0.024)	(0.020)	(0.019)	(0.023)	(0.023)		(0.015)	(0.019)
LHHLL	0.124	0.086	0.010	0.048	0.029		0.067	0.01
	(0.011)	(0.007)	(0.001)	(0.004)	(0.003)		(0.006)	(0.001)
LHHLH	0.105	0.036	0.019	0.067	0.038		0.095	0.143
	(0.009)	(0.003)	(0.002)	(0.006)	(0.004)		(0.008)	(0.012)
LHHHL	0.086	0.064	0.057	0.057	0.143		0.076	0.048
	(0.008)	(0.005)	(0.005)	(0.005)	(0.012)		(0.007)	(0.004)
LHLLL	0.048	0.021	0.067	0.000	0.000		0.057	0.010
	(0.004)	(0.002)	(0.006)	–	–		(0.005)	(0.001)
HHHHH						0.629		
						(0.023)		
All Other	0.086	0.05	0.029	0.124	0.371	0.000	0.200	0.114
	(0.008)	(0.004)	(0.003)	(0.011)	(0.023)		(0.016)	(0.010)

TABLE 13. Common Sequences of Actions (Last 5 supergames)

Treatment	5 or more observed Supergames				
En-DPD	CC,**DC**,**DD**,CC,**DD** 0.095	CC,**CC**,**CC**,**CC**,**CC** 0.086	CC,**CC**,**CC**,**CC**,**DC** 0.057	CC,**DC**,**DC**,**DC**,**DC** 0.057	
En-DPD-CC	CC,**CC**,**CC**,**CC**,**CC** 0.264	CC,**CC**,**CC**,**DC**,**DD** 0.079	CC,**CC**,**CC**,**CC**,**DC** 0.050	CC,**DC**,**DD**,CD,CD 0.043	
En-DPD-HT	CC,**CC**,**CC**,**CC**,**CC** 0.610	CC,**CC**,**CC**,**CC**,**DC** 0.124			
En-DPD-X	CC,**CC**,**CC**,**CC**,**CC** 0.162	CC,**CC**,**CC**,**CC**,**DC** 0.095	CC,**DC**,**DC**,**DC**,**DC** 0.048		
Ex-DPD	DD,DD,DD,DD,DD 0.381	DC,DD,DD,DD,DD 0.352			
Ex-SPD (Low)	CC,CC,CC,CC,DC 0.128	DC,DD,DD,DD,DD 0.128	CC,CC,CC,CC,CC 0.128		0.629
Ex-SPD (High)	**DD**,**DD**,**DD**,**DD**,**DD** 0.515	**DC**,**DC**,**DD**,**DD**,**DD** 0.091	**DC**,**DD**,**DD**,**DD**,**DD** 0.076		
En-DCP-M	CC,**DC**,**CD**,**DC**,**CD** 0.076	CC,**DC**,**DD**,CC,**DD** 0.067	CC,**DD**,CC,**DD**,CC 0.057	CC,**DD**,CC,**DD**,DC 0.057	DC,DC,DD,DD,DD 0.057
En-DCP-E	CC,**D**D,CC,**DD**,CC 0.257	CC,**DC**,**DD**,CC,**DD** 0.105	CC,**DC**,**CD**,**DC**,**CD** 0.095	CC,**DD**,CC,**DD**,DC 0.095	CC,**DC**,**DC**,**DC**,**DC** 0.048

Note: All treatments except for Ex-DPD display *High*-state action pairs in bold face.

Robustness of the SFEM estimates. The estimates reported in Table 14 result when the strategies included correspond to those that capture most behavior in infinitely repeated prisoner's dilemma experiments. For each treatment we include always cooperate (M_{CC}), always defect (M_{DD}), the best Markov perfect equilibrium whenever it differs from M_{DD}, a trigger strategy with reversion to the best Markov perfect equilibrium and Tit for Tat. Comparing the measure of goodness-of-fit (β) to the estimates in Table 4 we observe only a minor reduction. This suggests that this simple set of strategies can rationalize our data to a large extent.

For treatments where the efficient outcome can be supported with A_{DD} or A_{CD} Table 15 reports the estimates using the two versions of each strategy depending on whether the strategy starts by selecting C or D the first period the game is at the high state (for more details in footnote 21). In En-DPD the estimates remain largely unchanged except that the frequency of strategy that starts by cooperating and punishes with M_{CD} after a deviation, which we call A^C_{CD}, is above 20%. Comparing to the estimates in Table 4 we verify that there is a reduction of similar magnitude in the estimate of S_{CD}. This highlights the difficulty of identifying a strategy such as A^C_{CD} from S_{CD}: both strategies prescribe to cooperate in high if there are no previous deviations and would coincide from then on if there is no coordination on alternation in the second period in high. Other than these discrepancies the estimates reported in Table 4 remain largely unchanged.

Table 16 presents estimates when we expand the set of Markov strategies in En-DPD-X, where we change the size of the state-space. To explain the extra strategies, consider first Figure 4. The figure presents the cooperation rates in low and in high in panels (A) and (B), respectively. Supergames are grouped in blocks of five and the state-space X is divided in three parts: lower than or equal to -3, between -3 and 3, and higher than or equal to 3. Panel (A) shows that the cooperation rate in low is largely unaffected by the choice of x. However, for high state in panel (B) there is a positive effect on cooperation as values of x are higher. Guided by this figure we included two extra strategies in our estimation $M^x_{CCC,DCC}$ and $M^x_{CCC,DDC}$. The supra-script indicates that it is a Markov strategy that conditions on x. The first (last) three values of the subindex indicate the action prescribed in the low (high) state for each of the three elements in the partition of X. Both strategies prescribe the choice of C in the low state for all values of x. This is consistent with the high cooperation rates in panel (A) of Figure 4. In the high state, strategy $M^x_{CCC,DCC}$ prescribes to defect only if the value of x is lower than or equal to -3, while $M^x_{CCC,DDC}$ would also defect if x is between -3 and 3. We also include trigger strategies that aim to implement joint cooperation, but use either of these strategies as punishments ($S^x_{CCC,DCC}$, $S^x_{CCC,DDC}$).

The estimates in Table 16 are significant a only in the case of $M^x_{CCC,DCC}$, reaching approximately one-fifth of the mass. Relative to the estimates in Table 4, the reduction is coming from M_{CC} and S_{CD}. The inclusion of these strategies, however, only leads to a minor improvement in the measure of goodness-of-fit, from 0.828 to 0.846.

Table 18 presents a different type of robustness exercise, in this case by reporting histories that are perfect fit to some strategy. The SFEM allows subjects to implement a strategy with errors, but as a benchmark we also report by treatment the proportion of histories that fit some strategy without errors. We start by commenting the last two rows of the table, that report the proportion of histories by treatment that are a perfect fit to at least one of the identified strategies (explained histories), and the complement (unexplained histories). Overall, approximately 58% of all histories are explained, which indicates that the identified set of strategies -a set that is small relative to all possible strategies- can be quite successful in perfectly describing behavior in a majority of histories. The proportion of explained histories does vary by treatment, with the lowest figure at slightly above than 40% and the highest being close to 85%.

To report figures of perfect fits by strategy and treatment we proceed in the following manner. For each treatment we count the number of histories that can be perfectly fit by each strategy. In some cases it is possible that the same history can be perfectly fit by more than one strategy. We then compute the total number of perfect fits by treatment (which may include a history more than once), and report in the first two blocks of rows of Table 18 the proportion of perfect fits relative to the total. Comparing the proportions reported in the Table to the SFEM estimates in Table 4 shows that we would reach the same qualitative conclusions with either type of analysis.

TABLE 14. SFEM Output: Constrained Set of Strategies (Last 5 Supergames)

Strategies	**En-DPD**	**En-DPD-CC**	**En-DPD-HT**	**En-DPD-X**	**Ex-DPD**	**Ex-SPD** ($\theta = L$)	**Ex-SPD** ($\theta = H$)	**En-DCP-M**	**En-DCP-E**
Markov									
M_{CC}	0.118*	0.177**	0.206*	0.362***	0.000	0.104*	0.000	0.066	0.092
	(0.072)	(0.080)	(0.122)	(0.092)	(0.015)	(0.060)	(0.014)	(0.053)	(0.059)
M_{DD}	0.045	0.087	0.032	0.027	0.630***	0.268***	0.733***	0.265**	0.069
	(0.042)	(0.058)	(0.041)	(0.041)	(0.137)	(0.092)	(0.118)	(0.122)	(0.054)
M_{CD}	0.265	0.115	0.050	0.189	0.093**			0.393	0.682
	(0.102)	(0.060)	(0.033)	(0.071)	(0.047)			(0.092)	(0.109)
History-dependent									
S_{DD}					0.194	0.168	0.080		
					(0.109)	(0.100)	(0.061)		
S_{CD}	0.207**	0.194**	0.041	0.301**				0.077	0.088*
	(0.082)	(0.080)	(0.078)	(0.122)				(0.062)	(0.052)
TfT	0.365	0.427	0.671	0.121	0.083	0.460	0.187	0.198	0.069
β	0.813	0.852	0.920	0.819	0.943	0.855	0.888	0.776	0.862
# Observations	1,260	1,568	1,358	1,204	1,050	454	764	1,232	1,218

Note: Bootstrapped standard errors in parentheses. Level of Significance: ***-1 percent; **-5 percent; *-10 percent.

TABLE 15. SFEM Output including both versions of A_{CC} and A_{CD} (Last 5 Supergames)

Strategies	**En-DPD**	**En-DPD-CC**	**En-DPD-X**	**En-DCP-M**	**En-DCP-E**
			Markov		
M_{CC}	0.117*	0.173**	0.347***	0.068	0.092
	(0.071)	(0.077)	(0.090)	(0.052)	(0.058)
M_{DD}	0.024	0.039	0.027	0.077	0.048
	(0.030)	(0.039)	(0.034)	(0.061)	(0.038)
M_{CD}	0.211*	0.057	0.138*	0.311**	0.651***
	(0.127)	(0.063)	(0.081)	(0.157)	(0.180)
M_{DC}	0.000	0.000	0.000	0.063	0.000
	(0.003)	(0.010)	(0.014)	(0.042)	(0.001)
			History-dependent		
S_{DD}	0.106	0.227*	0.069	0.000	0.000
	(0.102)	(0.124)	(0.065)	(0.042)	(0.008)
S_{CD}	0.000	0.075	0.245**	0.000	0.088
	(0.076)	(0.069)	(0.123)	(0.052)	(0.056)
TfT	0.232***	0.304***	0.089	0.134*	0.069
	(0.083)	(0.094)	(0.061)	(0.079)	(0.078)
$STfT$	0.023	0.016	0.021	0.000	0.000
	(0.027)	(0.020)	(0.030)	(0.009)	(0.002)
A^D_{DD}	0.059	0.046	0.029	0.161	0.051
	(0.039)	(0.041)	(0.035)	(0.109)	(0.050)
A^D_{CD}	0.000	0.065	0.035	0.080	0.000
	(0.103)	(0.060)	(0.054)	(0.130)	(0.143)
A^C_{DD}	0.000	0.000	0.000	0.044	0.000
	(0.058)	(0.076)	(0.055)	(0.038)	(0.002)
A^C_{CD}	0.228	0.000	0.000	0.061	0.000
β	0.826	0.869	0.828	0.806	0.870
# Observations	1,260	1,568	1,204	1,232	1,218

Note: Bootstrapped standard errors in parentheses. Level of Significance: ***-1 percent; **-5 percent; *-10 percent.

TABLE 16. SFEM Output: Additional Strategies in En-DPD-X

Strategies	**En-DPD-X**
	Markov
M_{CC} (M_{CCCD})	0.253***
	(0.078)
M_{DD} (M_{CDDD})	0.027
	(0.034)
M_{CD} (M_{CCDD})	0.133*
	(0.071)
M_{DC} (M_{CDCD})	0.000
	(0.013)
M_{CCCC}	
M_{DDDD}	
$M^{x}_{CCC,DCC}$	0.203**
	(0.098)
$M^{x}_{CCC,DDC}$	0.002
	(0.048)
	History-dependent
S_{DD} (S_{CDDD})	0.073
	(0.062)
S_{CD} (S_{CCDD})	0.162
	(0.119)
$S^{x}_{CCC,DCC}$	0.000
	(0.019)
$S^{x}_{CCC,DDC}$	0.000
	(0.020)
TfT	0.063
	(0.056)
$sTfT$	0.015
	(0.024)
A_{DD} (A_{CDDD})	0.032
	(0.036)
A_{CD} (A_{CCDD})	0.038
β	0.846
# of Observations	1,204

Note: Bootstrapped standard errors in parentheses. Level of Significance: ***-1 percent; **-5 percent; *-10 percent.

TABLE 17. SFEM Output: Strategies of the Selection Index (Last 5 Supergames)

Strategies	**En-DPD**	**En-DPD-CC**	**En-DPD-HT**	**En-DPD-X**	**Ex-DPD**	**Ex-SPD** ($\theta = L$)	**Ex-SPD** ($\theta = H$)	**En-DCP-M**	**En-DCP-E**
M_{CD}	0.328	0.210	0.064	0.145	-	-	-	0.740	0.757
M_{DD}	-	-	-	-	0.662	0.370	0.774	-	-
S_{CD}	0.672	0.790	0.936	0.855	-	-	-	0.260	0.243
S_{DD}	-	-	-	-	0.338	0.630	0.226	-	-
β	0.720	0.753	0.845	0.710	0.927	0.739	0.874	0.607	0.796
# Observations	1,260	1,568	1,358	1,204	1,050	454	764	1,232	1,218
β (Table 4)	0.826	0.869	0.940	0.828	0.947	0.868	0.902	0.805	0.870
Difference β (Table 4)-β	0.106	0.116	0.095	0.118	0.020	0.129	0.028	0.198	0.097

Note: For each treatment the estimation includes only the strategies used in the index presented in Section 6.2. The median (mean) of difference between β in this estimation and that of Table 4 is 0.116 (0.981).

TABLE 18. Proportion of histories that are perfect fits

Strategies	**En-DPD**	**En-DPD-CC**	**En-DPD-HT**	**En-DPD-X**	**Ex-DPD**	**Ex-SPD** ($\theta = L$)	**Ex-SPD** ($\theta = H$)	**En-DCP-M**	**En-DCP-E**
Markov									
M_{CC}	0.234	0.216	0.241	0.232	0.028	0.217	0.027	0.041	0.097
M_{DD}	0.000	0.022	0.007	0.020	0.304	0.169	0.460	0.159	0.025
M_{CD}	0.082	0.049	0.009	0.052	0.073			0.304	0.609
M_{DC}	0.000	0.022	0.007	0.020	0.014			0.159	0.025
History-dependent									
S_{DD}	0.221	0.227	0.243	0.210	0.129	0.244	0.076	0.065	0.040
S_{CD}	0.227	0.219	0.243	0.230	0.037			0.067	0.099
TfT	0.220	0.225	0.236	0.218	0.109	0.268	0.074	0.099	0.069
$STfT$	0.008	0.004	0.002	0.007	0.212	0.102	0.363	0.041	0.002
A_{DD}	0.004	0.008	0.005	0.006	0.084			0.032	0.017
A_{CD}	0.004	0.008	0.005	0.006	0.010			0.032	0.017
Explained histories	0.500	0.551	0.679	0.486	0.840	0.644	0.773	0.427	0.525
Unexplained histories	0.500	0.449	0.321	0.514	0.160	0.356	0.227	0.573	0.475
Total # of histories	630	840	630	630	630	250	380	630	630

Note: For each treatment each cell in the Markov and History-dependent blocks reports the proportion of histories that are a perfect fit to the corresponding strategy relative to all perfect fits. The last two rows report the proportion of explained and unexplained histories by treatment. A history is considered unexplained if it is not a perfect fit to any of the identified strategies.

Appendix B. For Online Publication: Markov Quantal Response Equilibrium

An alternative hypothesis to the differential selection over history-dependent and state-dependent equilibria that we posit in the main paper is that the observed deviations are caused by noisy Markov Perfect play. Under this hypothesis, subjects play a noisy best response, given the state and the other player's mixed strategy. We now briefly summarize the theory for this (drawing heavily from Goeree et al. (2016))

Given a state dependent mixed strategy over the available state actions (a probability distribution σ_i^θ over actions at each state θ, for each player i) we can solve the following linear-system to get expected value at each state:

$$V^\star(\theta;\sigma) = \mathbb{E}_\sigma\left[(1-\delta)\cdot u_i\left(\sigma^\theta(a),\theta\right) + \delta\cdot V^\star\left(\psi\left(\sigma^\theta(a),\theta\right);\sigma\right)\right].$$

Given this, we can calculate the expected value of each specific action choice $a_i \in \mathcal{A}_i$ as

$$V_a^\star(\theta;\sigma) = \mathbb{E}_\sigma\left[(1-\delta)\cdot u_i\left(\left(a,\sigma_{-i}^\theta(a)\right),\theta\right) + \delta\cdot V^\star\left(\psi\left(\left(a,\sigma_{-i}^\theta(a)\right),\theta\right);\sigma\right)\right].$$

A logit Markov-perfect quantal response equilibrium (logit-MQRE) is defined as a mixture that solves the series of fixed points

$$\tilde{\sigma}^\theta(a) = \frac{e^{\lambda\cdot V_a^\star(\theta;\tilde{\sigma}^\theta)}}{\sum_{b\in\mathcal{A}_i} e^{\lambda\cdot V_b^\star(\theta;\tilde{\sigma}^\theta)}},$$

for all states $\theta \in \Theta$ and actions $a \in \mathcal{A}_i$, where the parameter $\lambda \geq 0$ controls the degree of noise. When $\lambda = 0$ the game is pure noise. As $\lambda \to \infty$ the game tends to the standard mutual-best-response restriction of an MPE.

The computations are presented in Figure 5 indicating logit-MQRE predictions as we trace out the locus of state-dependent cooperation rates $\tilde{\sigma}(\lambda) = \left(\tilde{\sigma}^L(C;\lambda), \tilde{\sigma}^H(C;\lambda)\right)$ in all of our two-state games shifting λ from close to zero (the point $(1/2, 1/2)$ in all examples) through to very large values (symmetric MPEs of each game).[46]

Alongside these theoretical predictions we provide the sample state-dependent average cooperation rates in the last five cycles of each treatment $\hat{\sigma} = (\hat{\sigma}^L(C), \hat{\sigma}^H(C))$ as a gray circle on each diagram. Examining the position of the gray circle relative to the logit-MQRE locus in most treatments the first conclusion is that the MQRE predictions seems to succeed in the majority of treatments. Except for En-DPD-HT, the sample average $\hat{\sigma}$ is never too far away from some point on the MQRE prediction locus $\tilde{\sigma}(\lambda)$.

[46] In some cases there are multiple-equilibria, and we attempt to illustrate all such outcomes, though we do not illustrate the exact value of λ below which some equilibria cease to be fixed points.

From another point of view, however, in some treatments (En-DPD, En-DPD-CC, En-DPD-HT) the noise parameter λ that best fits the sample averages is close to zero. In other words, the logit-MQRE that best fits the data involves the largest possible amount of noise. Moreover, these are treatments where the data exhibits substantial variation if we divide the cooperation rates into two cases: cooperation by state conditional on no defections in a prior round (white square in the diagram), and cooperation by state conditional on a defection in a prior round (black square). This therefore divides the data into the two history-dependent cases that a grim-trigger strategy would use. While the resulting average (gray circle) can be close to a logit-MQRE with low λ (e.g. En-DPD), the decomposition suggests the presence of a consistent history dependent response more than play of a MPE with a relatively large level of noise.

Contrarily, in En-DCP treatments the decomposition (white square/black square) is relatively closer to the mean (gray circle) and the λ that best fits the data is relatively higher -lower level of noise needed to rationalize the data. This suggests that the logit-MQRE is a better predictor of behavior in these treatments. A similar finding holds for Ex-DPD.[47]

Overall, our findings here are consistent with the main reports in the paper. In treatments where Markov strategies can better rationalize subjects' choices (e.g. En-DCP), the logit-MQRE is a better fit. In treatments where history-dependent strategies rationalize a large proportion of the data (e.g. En-DPD-HT), the logit-MQRE does not provide a good fit.

[47] While in Ex-DPD the average for cooperation rates conditional on no defections in prior rounds (white square) is far from the other averages it is worth noting that there are relatively few instances of prior rounds with no defections.

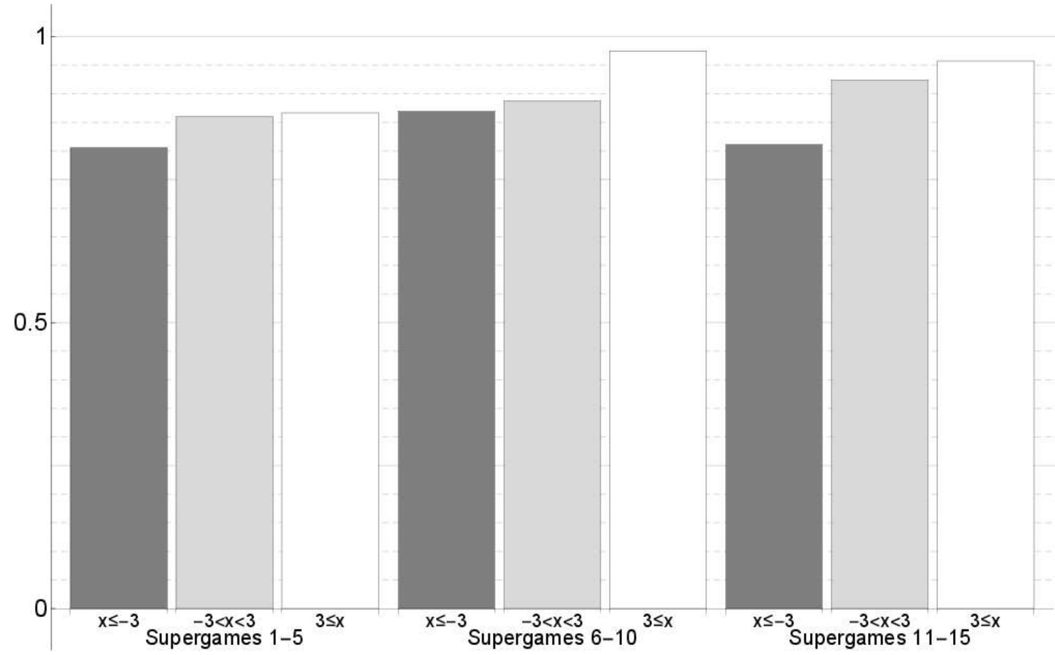

(A) Low state, period one

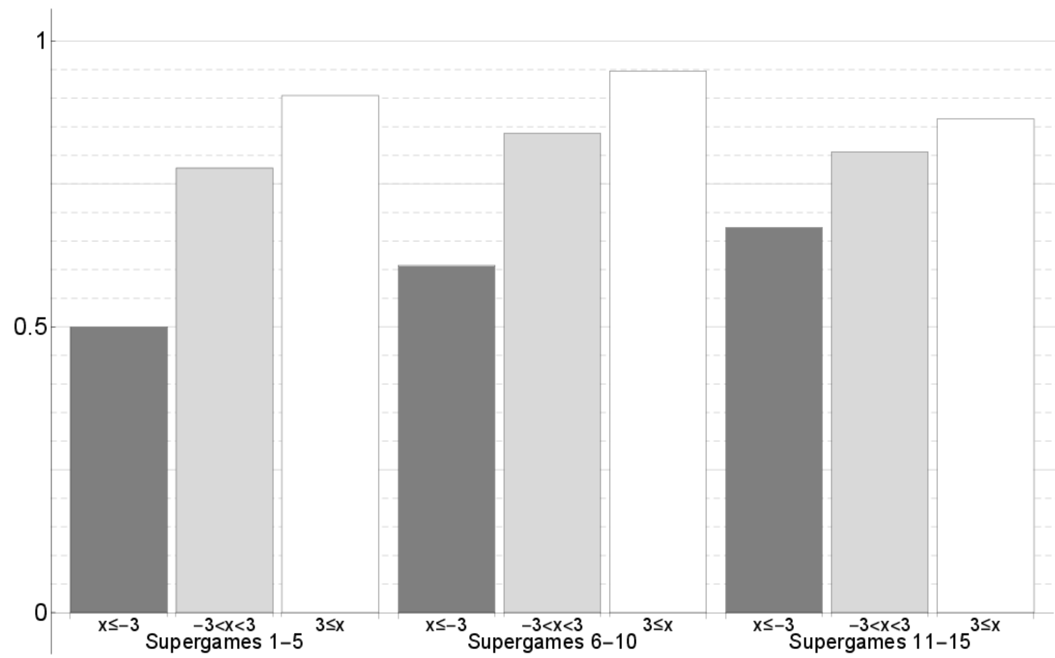

(B) High state, period two

FIGURE 4. Cooperation rates in En-DPD-X

Note: The unit of observation is a period. Running a random-effects probit estimates, for the low state in period one, only the difference between cooperation for $x \leq -3$ and $x \geq 3$ is significant (95 percent confidence, for both supergames 6–10 and for 11–15). For the high-state cooperation in period two, the difference between cooperation for $x \leq -3$ and $x \geq 3$ is always significantly different (above 99 percent confidence, each block of five).

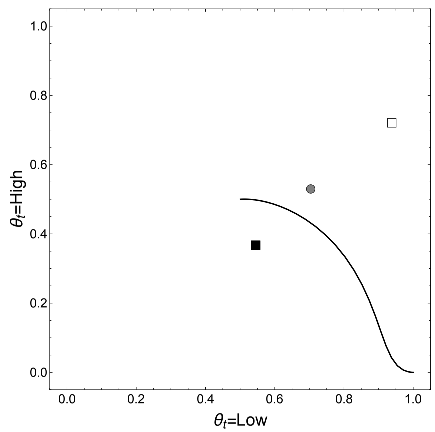

(A) En-DPD

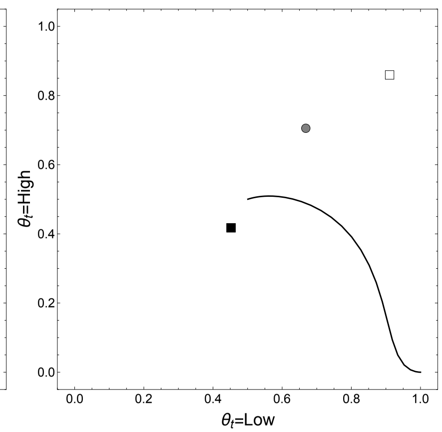

(B) En-DPD-CC

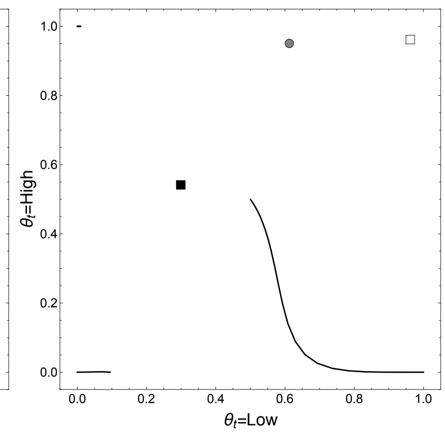

(C) En-DPD-HT

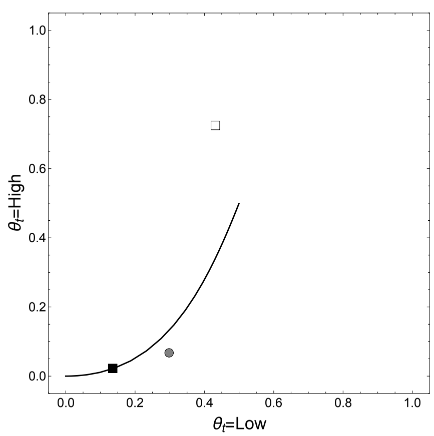

(D) Ex-DPD

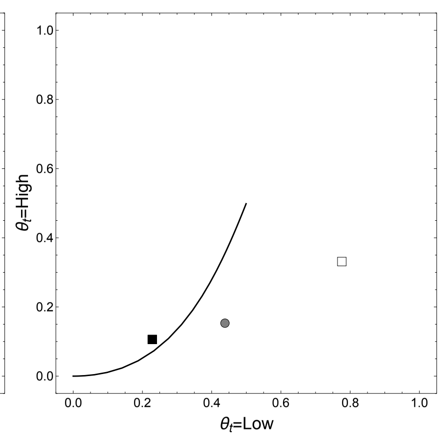

(E) Ex-SPD

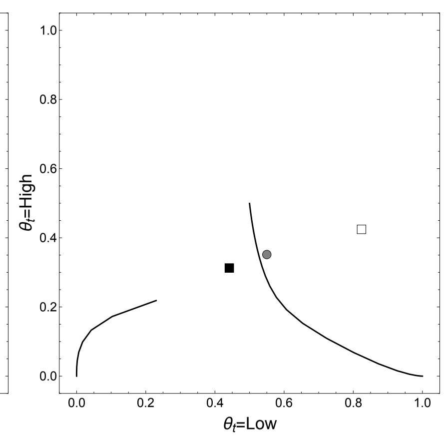

(F) En-DCP-M

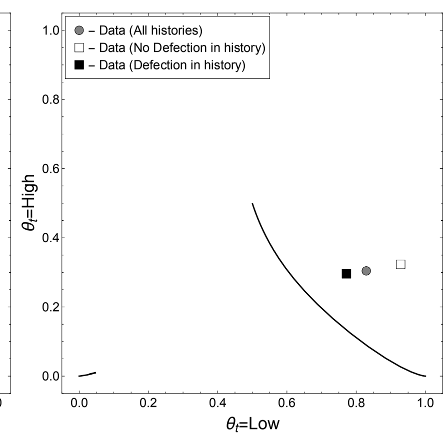

(G) En-DCP-E

FIGURE 5. Markov QRE

Note: Black lines indicate the proportion of cooperation at each state in a logit Markov Quantal Response Equilibrium, for some value of the noise parameter λ. Three points represent: Grey circle–State-conditioned average in the last five cycles; White square–state-conditioned average for all cycles where no defection has been observed in a previous round; Black square–state-conditioned average for cycles where a defection was observed in a previous round. Note that as the state is endogenous and conditioned here, the circle will not necessarily be a convex combination of the other two points.