DYNAMIC INCENTIVES AND MARKOV PERFECTION: PUTTING THE 'CONDITIONAL' IN CONDITIONAL COOPERATION

EMANUEL VESPA AND ALISTAIR J. WILSON

ABSTRACT. Many economic applications, across an array of fields, use dynamic games to study strategic interactions that are dynamic in nature. While these games will generically have large sets of possible equilibria, Markov perfection (MPE) is the main criterion for selection in applied work. Our paper experimentally examines this assumed selection across a number of simple dynamic games. Starting from a two-state modification of the most studied static environment–the infinitely repeated PD game—we work outward, characterizing the response to broad qualitative changes to the game's features. Subjects in our experiments show an affinity for conditional cooperation, readily conditioning their behavior not only on the state but also the recent history of play. More-efficient history-dependent play is the norm in many treatments, though the frequency of MPE-like play can be predicted with a modification to an index developed for infinitely repeated games. A dynamic extension of the basin of attraction is shown to have predictive power for the selection of MPE outcomes.

1. INTRODUCTION

The trade-off between opportunistic behavior and cooperation is a central economic tension. In settings where agents interact indefinitely theory shows it is possible to support efficient outcomes. So long as all parties place enough weight on long-run benefits from sustained cooperation, threats to condition future behavior on the present outcome are credible, and powerful enough to deter opportunistic choices. This holds whether the strategic environment is fixed (a repeated game) or evolving through time (a dynamic game). The set of subgame-perfect equilibria (SPE) is large, with many equilibrium outcomes possible across a range of efficiency levels. For repeated games—a special case within dynamic games—the experimental literature has documented a number of patterns for behavior (see Dal Bó and Fréchette, 2014, for a survey). In comparison, much less is known for the larger family of dynamic games. In this paper we expand outward from what is already well-known, experimentally investigating how behavior in very simple dynamic games responds to broadly read features of the environment.

Date: July, 2015.
We would like to thank Gary Charness, Pedro Dal Bó, John Duffy, Ignacio Esponda, Guillaume Fréchette, Drew Fudenberg, Ryan Oprea and Lise Vesterlund, as well as seminar audiences at Brown, Caltech, Lafayette College, Michigan, Pittsburgh, UC Santa Barbara, UC Riverside, the 2014 ESA meetings, the 2015 Conference on Social Dilemmas, and the Jerusalem theory summer school.

Dynamic games are frequently used in both theoretical and empirical applications, and the analysis typically requires some criterion for equilibrium selection.[1] In principle, just as with repeated games, strategies can condition on the observed history of play. Such history-dependent strategies can bootstrap cooperative outcomes in equilibrium, for example through trigger strategies that cooperate conditional on no observed deviations, otherwise switching to an incentive-compatible punishment phase. Yet the most-common solution concept in the dynamic-games literature precludes such history-dependence. Instead, the literature focuses the search for equilibria on those strategies where agents condition their choices only on the present "state" of the game—where each state in the dynamic game corresponds to a specific stage-game.[2]

Strategies that condition the selected action *only* on the present state are referred to as Markov strategies. While analytically tractable, because Markov strategies are memoryless they cannot punish based on observed deviations from the intended path of play. Typically, strategies that condition on the larger history can sustain efficient outcomes in equilibrium, where Markov strategies with their tighter conditioning can not. Where the emphasis in repeated games is on equilibria that use past play to support efficient outcomes, the focus on Markov in more general dynamic games ignores such conditioning—potentially ruling out efficient outcomes, even those supportable through an SPE.

The available experimental evidence on behavior mirrors this rift. On the one side, a large experimental literature on the infinite-horizon prisoner's dilemma (PD) game documents a majority of subjects using efficient, history-dependent strategies when the future discount rate is large enough that these strategies are equilibria. On the other side, a nascent literature on infinite-horizon dynamic games suggests that behavior is consistent with the subset of SPEs where players do use Markov strategies (Markov perfect equilibria, MPE).[3] Our paper's aim is to connect the experimental literatures on infinitely repeated and dynamic games, and provide evidence to predict situations where the selection of MPE is more likely, and in which it is unlikely. A natural task then is characterizing which properties of a dynamic game might lead to the selection of state-dependent

[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).

[2]Here we refer to the notion of Markov states, which are endogenously defined as a partition of the space of histories (see Maskin and Tirole 2001 for details). 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), page 178.

[3]Battaglini et al. (2012) were the first to provide experimental evidence where the comparative statics are well organized by MPE. See Battaglini et al. (2014) and Salz and Vespa (2015) for further evidence. In Vespa (2015), the choices of a majority of subjects can be rationalized using Markov strategies. For other experiments with infinite-horizon dynamic games see Saijo et al. (2014) and Kloosterman (2015).

behavior, and which to less-restrictive history dependence. Hence our paper's subtitle: what is the "conditional" in conditional cooperation, states or actions?

The experimental literature on dynamic games has primarily focused on rich dynamic environments, with many possible states. At the other extreme, the infinitely repeated PD is effectively a degenerate dynamic game, with just a single state variable (and an MPE of joint-defection forever). One simple characterization might be that behavior becomes Markovian as soon as the state-space is non-degenerate. We will show this is not the case. In our core two-state environment, more-efficient SPE are the norm, where the path of play after miscoordinations identifies history-dependent behavior. After showing this simple one/many distinction does not work to predict the majority of behavior and outcomes, we move on to isolate and modify other qualitative features of our core game.

This core game—which we will call our "pivot"—extends the most-studied indefinitely repeated game by adding a single additional state. In both of these states agents face a PD stage game. However, the payoffs in the *Low* state are unambiguously worse than those in the *High* state, where the best payoff in *Low* is smaller than the worst payoff in *High*. 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 has 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 strategies. Our modifications to this pivot involve *eight* between-subject treatments, where each examines how a change to an isolated feature of the original game affects behavior and the selection of strategies.

In the pivot (and many of our modified versions of it) we find that a majority of subjects seek to support efficient outcomes with history-dependent choices, at comparable levels to those reported for infinitely repeated PD games. This indicates a smoother transition over selection from infinitely repeated to dynamic games—the mere presence of an additional state does not drive subjects to ignore history and focus solely on the state. This is not say that Markov play is non-existent in our data, and importantly, where we do observe it, it is consistent with theory. 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 set of manipulations alters the efficient frontier in the dynamic game, making a single symmetric SPE more focal while holding constant the MPE prediction. Weakening the temptation to defect in the high state, we make coordination on the best-case SPE easier, and our treatments assess the degree to which behavior responds away from the MPE. A "static" manipulation alters a single payoff at a single state (reducing 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 pivot's stage-game payoffs, we make it harder to remain in *High*). In both manipulations the direction of the change in behavior is an increase in the selection of efficient outcomes, with equilibrium Markov play becoming negligible. Quantitatively, the effect is much stronger in the dynamic manipulation, where efficient play represents approximately three-quarters of the observed data.

The second set of manipulations focuses on how one agent's chosen action affects the other participant, on the nature of the strategic externalities. In our pivot one subject's current choices can affect both the other's current payoffs (a static externality) and the other's future payoff (a dynamic externality operating through the state's transition). To what extent are each helping to support history-dependent play?

We remove the pivot's dynamic externality in two distinct ways. In the first, we make the transition between the two states exogenous, but where both states are reached (stochastically) along the game's path. In the second, we remove the dynamics entirely, playing out each of the pivot's two stage-games as separate infinitely repeated games. In both manipulations, the only MPE involves playing the stage-game Nash: unconditional joint defection. Relative to the pivot, we observe substantially less cooperation in both treatments—thus, dynamic externalities are shown to be an important selection factor for the supra-MPE behavior in the pivot. To remove static externalities we require that each agent's contemporaneous choice does not affect the other's contemporaneous payoff. We conduct two separate parametrizations, in which the broad structure of the equilibrium set remains comparable to the pivot: the efficient actions are exactly the same (and stay in the *High* state), while the most-efficient MPE still alternates between the *Low* and *High* states. In both parametrizations we again 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.

The final set of manipulations involves increases to the pivot game's state-space. One argument often informally made in favor of the Markov restriction is that when environments are "complex," agents may find it easier to use "simple" strategies. Here our manipulations increase the number of possible states, while holding constant many elements from the pivot, to explore whether there is a resulting increase in Markov play. Our first state-space manipulation is a perturbation, adding small-scale, exogenous noise to the pivot's payoffs. Each shock to the game is an independent draw and its effect on the game is non-persistent, where only the *Low*/*High* component is endogenous. Our findings for this treatment indicate that despite an order of magnitude increase in the state-space's size, behavior is similar to the pivot. If anything we observe an increase in the frequency of cooperative play.

Our second state-space manipulation adds two endogenous states (each with their own particular stage games) to the pivot. We term these states *Very Low* and *Very High,* and choose their stage games so that the main structure of the pivot remains: the efficient and MPE outcomes do not change. We still find substantial use of history-dependent strategies. However, the treatment offers a new rationale for why Markov behavior may emerge: strategic uncertainty. With more endogenous states, there is greater variation in subjects' play, and it takes longer for this uncertainty to be realized. Our data reflects increased pessimism about others in this treatment, with more subjects choosing to start the game by defecting. Coordinating on efficient cooperative outcomes becomes more challenging, and greater rates of miscoordination lead to paths of play more consistent with state-dependent equilibria of the game.

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, where many of the non-Markov strategies we observe aim for efficiency. ii) For those subjects who do use Markov profiles, the MPE is focal. iii) Weakening the temptations to defect from efficient, history-dependent play increase the selection of supra-MPE outcomes. iv) the presence of both static and dynamic externalities affect coordination over history-dependent strategies, where removing either type of strategic externality leads to a much greater selection of MPE behavior. v) Increased complexity in the state-space does not on its own lead to MPE play becoming focal.

Clearly, the larger family of dynamic games is very rich, and our paper only looks at a small family of games within it. Our aim is to begin documenting which broad features of the environment have strong effects on behavior, so that eventually it might be possible to develop more-refined criteria for equilibrium selection in dynamic games, as has happened within the larger repeated-games literature. In the discussion section of the paper we expand on this, outlining some implications of our findings for this larger research agenda. In particular, we show that a simple dynamic extension of the basin of attraction used in the repeated-games literature has predictive power for the selection of more-efficient outcomes than the MPE.

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,... . Underlying the game is a payoff-relevant state $\theta_t \in \Theta$ evolving according to a commonly known transition rule $\psi : \mathcal{A} \times \Theta \to \Delta\Theta$, so that the state next round is given by $\theta_{t+1} = \psi(a_t, \theta_t)$. 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):

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

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 (possibly stochastic) 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 3–6 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

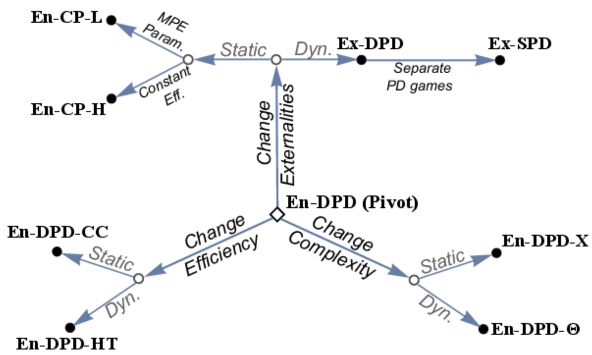

FIGURE 1. Summary of Treatment Design

Z_t is revealed to subjects directly after the decisions have been made for period t.[4] This method implements the expected payoffs in (1) under risk neutrality.[5]

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.[6] 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.[7]

2.4. **Overview of the design.** In total we will document results from nine distinct treatments, across three broad categories of manipulation: i) changing the efficient outcome to weaken the temptation to defect from history-dependent cooperation (Section 4); ii) changing strategic externalities, how one agent's choice affects the other's payoffs (Section 5); and iii) changing the size

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

[5]For payment we randomly select four of the fifteen supergames. 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."

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

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

TABLE 1. Treatment Summary

Treatment	$\lvert\Theta\rvert$	MPE		Efficient		IR action		Transition		$\Pr\{\theta_1 = L\}$
		L	H	L	H	L	H	L	H	
Pivot (Section 3):										
En-DPD	2	C	D	(C,C)	(C,D)	D	C	(C,C)	(D,D)	1
Change Efficiency (Section 4):										
En-DPD-CC	=	=	=	=	(C,C)	=	=	=	=	=
En-DPD-HT	=	=	=	=	(C,C)	D	D	=	not (C,C)	=
Change Strategic Externalities (Section 5):										
Ex-DPD	=	D	D	=	=	D	D	prob. 0.6	prob. 0.2	=
Ex-SPD	1	D	D	=	=	D	D	$\emptyset$	$\emptyset$	prob. 0.4
En-DCP-M	=	=	=	=	=	=	=	=	=	=
En-DCP-E	=	=	=	=	=	D	D	=	=	=
Change State-Space Complexity (Section 6:)										
En-DPD-X	22	$\doteqdot$	$\doteqdot$	$\doteqdot$	$\doteqdot$	=	=	$\doteqdot$	$\doteqdot$	$\doteqdot$
En-DPD-$\tilde{\Theta}$	4	$\doteqdot$	$\doteqdot$	$\doteqdot$	$\doteqdot$	=	=	$\doteqdot$	$\doteqdot$	=

Note: Where the table lists "=", the relevant cell is identical to the En-DPD game's value. For the En-DPD-X and En-DPD-$\tilde{\Theta}$ treatments 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.

of the state-space (Section 6). 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 Figure 1 and Table 1.

Figure 1 shows how all nine treatments are organized around the pivot (labeled En-DPD), while Table 1 summarizes the main differences in theoretical properties for each treatment, relative to the pivot. The table provides: i) the size of the state-space; ii) the most-efficient symmetric MPE; iii) the efficient action profile; iv) the action that obtains the individually rational payoff (by state); v) the action profile/probability of transition to a different state; and vi) the starting state θ_1. 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 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.[8] 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.

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

[8]An economic interpretation for this 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 a choice over that stock (extraction of fish or water, effluent from production, market supply). By cooperating in the low state, the stock can be built up to a socially-desirable level. Once at the high state, the stock is more robust and only transitions back to low following more systemic opportunistic behavior (joint defection).

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 stage-game Nash 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.[9]

From the point of view of identifying Markov behavior, we chose the pivot game 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} as a punishment.

[9] 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.

We will refer to this symmetric history-dependent trigger strategy with an M_{CD} punishment phase as S_{CD}.[10]

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.[11] 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 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.

Though 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 $\delta = 0.75$.[12,13] In particular, because all efficient outcomes can be supported as SPE, both players can receive discounted-average payoffs arbitrarily close to the first-best symmetric payoff of 178.75. As such, our pivot illustrates a tension not only between the best-case symmetric SPE and MPE, but also between what is achievable with symmetric and asymmetric strategies.

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.[14] The two figures are designed to illustrate aggregate-level behavior (Figure 3) and variation across supergames (Figure 4), while the table (Table 5)

[10]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 9 and implements an 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).

[11]We parametrize our pivot treatment with an asymmetric efficient outcome as this baseline will help when comparing with the manipulations of the strategic externalities in Section 5. Section 4 will present two treatments where symmetry is efficient; however the payoff difference between efficient and a symmetric solution is small, amounting to 5 cents per player.

[12]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).

[13]Every efficient supergame outcome $\{a_t\}_{t=1}^{\infty}$ in 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.

[14]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.

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 3, 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 3 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 show the cooperation rate in the last two periods in each supergame (again, conditional on the relevant state). For En-DPD, the illustrated 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.[15]

To further disaggregate behavior we move to Figure 4, 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.[16]

Figure 4(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

[15]Table 8 in the appendix provides the predicted 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.
[16]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.

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).[17] 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 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.

For the estimations reported in Table 5 we specify a very simple set of strategies.[18] It includes all four Markov strategies, 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.[19]

[17]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. .

[18]The SFEM output provides two inter-related goodness-of-fit estimates γ and β, and for comparability to other papers we report both. The parameter γ determines the probability of an error, and as $\gamma \to 0$ the probability that the choice prescribed by a strategy is equal to the actual choice goes to one. The probability that any choice is predicted correctly is given by the easier to parse β, which is a transformation of γ. Although the set of included strategies is simple, our measures of goodness-of-fit are far from a random draw (a β value of 0.5). This suggests that with this limited set of strategies it is possible to rationalize the data fairly well.

[19]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

The SFEM estimates for the pivot treatment, available in the first column of Table 5, reflect the heterogeneity observed in Figure 4(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.

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 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 here. 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

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 5 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}.

miscoordination. However, the broader results point to history-dependent play as the norm. The next three sections examine how modifications to the strategic environment alter this finding.

4. Changes to the Efficient Action

Our pivot is parametrized so the first-best outcomes are asymmetric. Our first set of treatments modify the pivot so the action maximizing the sum of the payoffs is unique and symmetric: joint cooperation. We achieve this through two distinct changes to the temptations to defect from joint cooperation: The first reduces the static temptation holding constant the continuation value from a defection. The second reduces the continuation value from a defection holding constant the static temptation.

4.1. **Static Change (En-DPD-CC).** Our first modification shifts the efficient actions by decreasing the payoff $u_i\left((D,C),H\right)$ from 280 to 250. All other features of the pivot—the starting state, the transition rule, all other payoffs—are held constant. 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 PD game: *En-DPD-CC*.

The data, presented in Figures 3 and 4(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.[20] Looking at the supergame level in Figure 4(B), this increase is reflected through larger concentrations in the top-left corner, perfectly cooperative supergames.

The estimated strategy weights in Table 5 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-punishment trigger S_{DD}.

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

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 matches this with an even weaker match to the MPE than the pivot. Our next treatment examines a similar exercise where we instead weaken the continuation value on a defection from joint-cooperation.

4.2. **Dynamic Change (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. Where 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 (cf. 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.

There are two broad changes relative to En-DPD from this shift in the dynamics: 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}.[21] Efficient joint cooperation is attainable as an SPE with either symmetric trigger, S_{DD} and S_{CD}.[22]

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

[21] 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.

[22] 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.

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 3 and comparing En-DPD-HT to the pivot, the most apparent results are the significant increase in high-state cooperation.[23] Comparing Figures 4(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 5 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.

While there is a clear increase in play that supports the efficient symmetric outcome, the SFEM estimates also indicates 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 best-case 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.[24]

4.3. **Conclusions.** In the two treatments above we reduce the payoff from a deviations from joint-cooperation in the high state. In a static treatment we alter this stage-game payoff from this deviation, and in a dynamic treatment we alter the continuation value. In both treatments the change makes symmetric cooperation the efficient outcome, and reduces the temptation to deviate from any history-dependent SPE that supports this outcome.

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

[24]The SFEM can identify two strategies that implement joint cooperation only if we observe some behavior in a punishment phase. Otherwise, two strategies such as S_{DD}, S_{CD} and M_{CC} are identical. Hence, when the procedure reports an estimate for M_{CC}, it can be capturing either M_{CC} or any history-dependent strategy that mostly cooperates and either does not enter its punishment phase within our data, or where that path is closer to M_{CC} than our other coarsely specified strategies. Vespa (2015) develops an experimental procedure to obtain extra information that allows to distinguish between such strategies and gain more identifying power.

Observed behavior in each treatment move towards efficient outcomes, and away from the MPE, though the effect is stronger in the dynamic manipulation (En-DPD-HT). One way to interpret these treatment effects is that the changes aid subjects' coordination, and reduce the effects of strategic uncertainty. Because our manipulations reduce the set of efficient outcomes (and SPE) the treatments 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.

5. Changes to the Externalities

In the above treatments there are two strategic considerations to 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 choice of action affects future values for the state, holding constant the En-DPD game's static externalities. 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.[25] Given this exogenous (Ex-) transition rule we label this dynamic PD treatment Ex-DPD.

[25]The Ex-DPD sessions were conducted after the En-DPD sessions were completed. The 60 percent and 80 percent probabilities were chosen to match aggregate state frequencies in the En-DPD sessions.

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 (given the strategy's history independence). 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.[26]

Looking at the experimental results for Ex-DPD, outcomes are starkly different from those where the state's evolution is endogenous. From Figure 3 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 3 is indicated at the supergame-level in Figure 4(D), 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 5. The highest frequency is attributed to the MPE, M_{DD}, with an estimate of just under 60 percent. For those subjects who 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.[27]

Removing the dynamic externalities dramatically shifts the observed behavior in the laboratory, leading to a collapse in cooperation. 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 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).

[26] 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.

[27] We also estimated strategy weights for this treatment adding the history-dependent strategy that supports cooperation only in the low state, described in footnote 26. 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.

Comparing aggregate-level results in Figure 3 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 4(E) 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.

SFEM estimates are presented by state in Table 5, 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 40 percent, substantially 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 three-quarters 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 stage-game across the session.[28]

Summary. Removing the dynamic externality from the pivot in Ex-DPD leads to a collapse of conditional cooperation, and the MPE becomes focal. However, when we remove the dynamics entirely, so that subjects face each stage game as a repeated game, we find an increases in the cooperation rate in both states relative to Ex-DPD. Having an evolving state within the supergame therefore makes it harder for subjects to cooperate. Equilibrium selection *does* respond to dynamic externalities, suggesting that the endogenously evolving state is a component in the selection of history-dependent cooperation in our pivot.

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.[29] Removing static externalities means the stage-game is no longer a PD game, and so we refer to this game instead as a dynamic common-pool (DCP) problem. For greater comparability with the pivot, two separate parametrizations are used, with stage-games presented in Table 3.

[28] In infinitely repeated PD, 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.

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

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.[30] 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. So 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}).[31]

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 3 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 4(F) shows a relatively large degree of variation across supergames, but with the largest mass concentrated at the bottom-right corner, consistent with the best-case MPE prediction, M_{CD}.

The SFEM estimates confirm the intuition from Figure 4(F), 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 best-case MPE. The difference in average cooperation rates between the two states is larger than in En-DCP-M (Figure 3), where the largest mass of supergames are in the bottom-right corner of Figure 4(G). 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.

[30]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.

[31]In addition, unlike the pivot not all efficient outcomes can be sustained as SPE for the DCP treatments.

Summary. Our dynamic common-pool treatments suggest increases in the selection of Markov strategies as we remove static externalities, in particular the best-case MPE M_{CD}. We do find some evidence for greater coordination on efficient asymmetric SPEs in En-DCP-M, relative to the pivot, where the treatment removes second-best symmetric SPEs (such as S_{CD}). 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. By comparing behavior in these treatments to the pivot we conclude that subjects' strategy selections do respond to the presence of static externalities, with an increase of symmetric MPE play as we remove them from the pivot.

5.3. **Conclusion.** Removing either the dynamic or static externalities in our En-DPD game weakens the theoretic desirability of more-efficient history-dependent strategies. For the treatments without endogenous dynamics, the power of history dependence is reduced as the future path of play can no longer be leveraged by the punishment path. In our treatments without static externalities, the change forces more-efficient outcomes to be asymmetric, and constraining to symmetric SPE, M_{CD} is the best outcome.

Presenting data from four treatments that are similar to the pivot but with each type of externality removed, we show that subjects' behavior responds with a greater selection of the relevant equilibrium Markov play than the pivot. The presence of both types of externality are therefore shown to be important to the selection of more-cooperative outcomes.

For the common-pool treatments it is possible that the absence of a more-efficient symmetric SPE is the primary driver for the increased Markov play, rather than the absence of static externalities. Though further research will likely separate between these forces more exactly, some evidence already exists. Vespa (2015) examines a dynamic common-pool game, but where the state-space has no upper bound, so that joint cooperation always leads to a higher payoff state. In his setting an efficient symmetric SPE exists, but modal behavior still mirrors the MPE prediction, suggesting that the absence of static externalities can be an independent driver for coordination on a state-conditioned response. His game also has a richer state-space, and this increased complexity may also contribute to the result. We turn to this selection channel in the next section.

6. Changes to the Complexity of the state-space

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, 2015). 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 two treatments examine this idea by manipulating the pivot game to increase 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.

The first of these treatments increases the number of payoff-relevant states from the pivot by adding exogenous, non-persistent shocks, independent of the original state variables. These shocks are moderately small in scale, and can therefore be thought of as a perturbation of the pivot's main strategic tensions, but with an order-of-magnitude increase in state complexity. The second treatment adds just two further states to the pivot—one below *Low*, the other above *High*—but both new states are associated with entirely distinct stage games, and are reached endogenously along the path of play. However, the two additional states are constructed to that neither the MPE nor the efficient outcomes should ever enter the new states.

6.1. **Static Complexity (En-DPD-X).** One simple way to add states while holding constant many of the pivot's strategic tensions is payoff-relevant noise. Our first complexity treatment adds a commonly known iid payoff shock each period through a uniform draw x_t over the support $X = \{-5, -4, \ldots, 4, 5\}$.[32] The specific payoffs in 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.

[32] This form of shock is common in IO applications that aim to structurally estimate the parameters of a dynamic game. See, for example, Ericson and Pakes (1995).

Moreover, the efficient frontier of the extended-state–space game is (for the most part) unaltered, as are the set of simple SPEs.[33] 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 3, we see qualitatively similar average cooperation rates to those in the pivot. Comparing Figures 4(A) and (H), this similarity extends to the supergame level, though the slightly greater cooperation in both states for En-DPD-*X* is a little more apparent.[34] To make the comparison across treatments cleaner, the SFEM estimates use the same strategies as our previous treatments, and thus ignore strategies that condition on the shock $x_t \in X$.[35] 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 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. Potentially the additional strategic uncertainty introduced into the game with the exogenous shock increases subjects' leniency.

Summary. Following our interpretation of this treatment as a perturbation of the pivot, the broad results point to a continuity in equilibrium selection with respect to the main strategic tensions of the dynamic game, where 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 their similar core strategic tensions.

6.2. **Dynamic Complexity (En-DPD-$\tilde{\Theta}$).** In the En-DPD-X treatment the additional state variable only affects the current period. Our final treatment enriches the state-space by adding two new states that can be endogenously reached along the path. In contrast to the first complexity treatment, our second alters the complexity over both what will happen this period (if the new states are reached) and conjectures over where the state is headed in future periods (at all states).

---

[33]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.

[34]By the last five rounds, the average behavior depicted in Figure 3 for En-DPD-*X* is significantly more cooperative in both states.

[35]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 5 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.

TABLE 4. En-DPD-$\tilde{\Theta}$ additional stage games

θ=Very Low (vL)

		2: C	2: D
1:	**C**	60,60	40,10
	D	10,40	20,20

θ=Very High (vH)

		2: C	2: D
1:	**C**	200,200	10,380
	D	380,10	85,85

Our dynamic-complexity treatment extends the state-space to $\tilde{\Theta} = \{vL, L, H, vH\}$ with the added states "very low" (vL) and "very high" (vH), where we label this treatment En-DPD-$\tilde{\Theta}$.[36]

The pivot's transition rule is extended to the four new states, with the same overall intuition: joint cooperation moves the supergame up a state (until the ceiling vH is reached); joint defection moves the supergame down a state (until the floor vH is reached). In all other cases the state is held constant. The new transition rule ensures all four states can be reached, and is given by

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

Payoffs in the low and high states are identical to those used in En-DPD, and we again start all supergames in the low state with certainty. For the two added states, the stage-game payoffs are given in Table 4.

The very-low stage game is chosen to have cooperation as the efficient, dominant strategy, where both players cooperating is both statically and dynamically efficient. We calibrate the payoffs so that when both players choose the dominant strategy of C in vL they receive the same payoff as joint defection in the pivot's low state. For our very-high state we choose a PD game (with D as the dominant strategy), and choose $u_i((C,C), vH) = 200$ so that symmetric joint cooperation in all periods yields an identical sequence of payments to the pivot. In order to keep the efficient outcomes identical to those in En-DPD, we set the off-diagonal payoffs in the vH stage-game so that alternation between (C,D) and (D,C) in this state is less efficient than the same alternation in H. However, we substantially increase the temptation to defect, with a payoff of 380. Finally, we set the joint defection payoff $u_i((D,D), vH)$ to \$0.85. This has the effect of making alternation

[36] In the experiment the additional states were referred to as the *Green Table* for vL and the *Orange Table* for vH.

between the high and low states superior to alternation between high and very high ($\frac{4}{7} \cdot 190 + \frac{3}{7} \cdot 130 > \frac{4}{7} \cdot 200 + \frac{3}{7} \cdot 85$).[37]

The pure-strategy MPE of the new game are (by construction) directly analogous to the pivot. The unique pure-strategy symmetric MPE is to cooperate in the low and very-low states, and defect in the high and very-high states. As such, along the path of play that starts in low, the supergame should only visit the low and high states under the MPE, yielding the same sequence of states and payoffs as M_{CD} in the pivot.[38] SPE exist that can attain any efficient outcome, as do symmetric SPE that maintain joint cooperation (in particular the analogues to S_{CD} and S_{DD}), yielding identical on-path payoffs to joint-cooperation in the pivot. However, our constructed game makes deviations from joint cooperation a dollar more tempting once vH has been reached.

Looking at the experimental results, overall cooperation rates are significantly lower than the pivot for both the low and high states (see Figure 3). The cooperation rates when the game reaches the vL state (which necessitates at least one joint-defection in low) are very high: 90 percent for the first five supergames falling to 85 percent for the last five. Cooperation rates in the vH state vary from 55 percent in the first five cycles to 60 percent in the last five. Surprisingly, the average cooperation rate in the vH state, with its more-powerful temptation to defect, is actually larger than in the H state.

At the history-level, the horizontal axis of Figure 4(I) measures the cooperation rate in either the vL or L states, and the vertical axis cooperation rates in either the H or vH states. Contrasting Figures 4(A) and (I) the observed patterns are similar, once we exclude the larger fraction of supergames that never get to the H-state or beyond (the supergames on the lower 'NaN' line).[39]

Comparing the pivot and En-DPD-$\tilde{\Theta}$ in the strategy estimates in Table 5 we observe similar total levels for Markov strategies, representing about one-third of the data in each case.[40] The equilibrium Markov strategy M_{CD} is still the most popular of the four, but the estimates do show an

[37]The individually rational payoff in the high state is still 130, while it is lower at $\frac{1}{4} \cdot 85 + \frac{3}{4} \cdot 130$ in the very high state. The individually rational payoff is 40 in the very low state, and is reduced from 60 to 45 in the low state.

[38]Paralleling En-DPD there is a pair of asymmetric MPE where one player defects in high and the other cooperates. The rest of the strategy is identical for the two players: cooperate in very low, defect in low, defect in very high. Given that low is the starting state, this MPE is inefficient, and alternates between low and very low. If this asymmetric Markov strategy is selected, there *is* a differing sequence of states selected relative to the sequence in En-DPD, however, the sequence of *payoffs* is constructed to be identical.

[39]For example, the cooperation rates in the high state when we exclude the supergames that do not reach the high state are similar. In the pivot, conditional on getting into the high state in period two (186 of 210 supergames), 38 percent manage to coordinate on joint cooperation. For En-DPD-Θ only 146 of the 210 supergames reach the high state, but 42 percent of these then coordinate on joint cooperation.

[40]For this treatment Table 5 and the text will abuse notation. A Markov strategy in this treatment should indicate the actions for each of the four possible states, so we could write the equilibrium strategy as M_{CCDD} for, respectively, cooperation in very low and low, and defection in high and very high. However, for simplicity and comparability to our other treatments we will restrict the Markov-strategies we look at to cooperate in very low and defect in very high. Where we show estimates for the strategy $M_{\sigma_L \sigma_H}$, we mean $M_{C\sigma_L \sigma_H D}$. Strategy estimation with more Markov

increase in M_{DD} and a decrease in M_{CC}, reflective of the increased frequency of supergames that never enter the high state. Since M_{CC} is the only Markov strategy profile which is not consistent with any equilibrium solution concept, but its outcome *is* consistent with successful history-dependent strategies, the estimates in En-DPD-$\tilde{\Theta}$ can be seen as pointing toward greater state-dependent equilibrium play.

In terms of history-dependent strategies the combined frequency does not change substantially from the pivot, but there are shifts over which particular strategies are used. Strategies supporting joint cooperation represent about a half of the estimates in En-DPD-$\tilde{\Theta}$, but the most common is the S_{DD} trigger, where this unforgiving trigger is a better response to the greater incidence of M_{DD}. In contrast, the pivot has TfT and S_{CD} as the two most-common cooperative strategies, and choosing M_{DD} is rare, both as an initial choice or through a trigger on miscoordination.

Relative to the pivot then, the main shift we observe is an increase in M_{DD}, both directly and indirectly as a punishment after a failed attempt at cooperation. While both players using M_{DD} is not sub-game perfect (either in En-DPD-$\tilde{\Theta}$ or the pivot) as there is a profitable deviation in the unreached high state, it is a Nash equilibrium. Moreover the induced path of play is consistent with an asymmetric MPE of the game $(M_{DD,}M_{DC})$.

In the pivot, the first-period cooperation rate is close to 95 percent so that less than one out of every ten supergames fails to enter the high state in period two. However, in En-DPD-$\tilde{\Theta}$ the initial cooperation rate is significantly lower (at approximately 80 percent), which leads to 36 percent of supergames failing to get to high in period two.[41] First round defections in the En-DPD-$\tilde{\Theta}$ treatment are particularly damaging: not only is the outcome that period inefficient, dynamically it becomes much more likely that the supergame will be stuck cycling between the low and very-low states. In period two, following an initial period where *one* player defects, only 12 percent of the 52 En-DPD-$\tilde{\Theta}$ supergames with this history manage to recoordinate on joint cooperation and reach the high state, while 46 percent are jointly defective and so move to vL. In contrast, 55 percent of supergames in the pivot with the same miscoordinated initial history (22 supergames) have joint cooperation in period two, and none have joint defection.

Subjects are therefore less likely to initially cooperate in En-DPD-$\tilde{\Theta}$ than the pivot, but also less willing to forgive these initial defections. Given the constructed similarities between the two

strategies for En-DPD-$\tilde{\Theta}$ is presented in Table 16 of the appendix, where we show that little is lost with the particular restrictions in vL and vH.

[41] A random-effects probit at the subject level assessed over the last five supergames rejects equivalence for first-period cooperation in En-DPD and En-DPD-$\tilde{\Theta}$ with 95 percent confidence.

games, the larger set of endogenous states does seem to make subjects less optimistic that the other will cooperate *in future states* and hence they are less likely to cooperate to begin with.[42]

6.3. **Conclusion.** The results in the complexity manipulations show that simply expanding the state-space on its own does not lead to a large increase of equilibrium Markov play. In fact, in En-DPD-X while we substantially increase the size of the state-space from 2 to 22 states, we find that cooperative outcomes are actually more likely. The presence of a large number of payoff-relevant states is not inhibiting subjects from coordinating on more-efficient outcomes than the MPE.

In En-DPD-$\tilde{\Theta}$ we have a smaller increase in the state-space (from two to four states), but where the new states are now endogenously reached. While we do find a small increase in equilibrium state-dependent strategies (from 23.5 percent in the pivot to 33.3 percent), the larger effects are greater variation over the strategies subjects are trying to coordinate over. Though there is not a huge increase in ex ante Markov play, greater miscoordination within the supergame leads to more paths of play that are eventually state-dependent and consistent with the MPE predictions.[43]

7. Discussion

7.1. **Summary of Main Results.** Our paper presents experimental results over a core pivot game, and eight modifications to it that create variation across three themes: i) coordination and efficiency; ii) the presence of different types of strategic externalities; and iii) the complexity of the state-space. Within each treatment we manipulate whether the changes are to the static or dynamic tensions within the game. We now summarize our main experimental results:

Result 1 (History Dependence)**.** *Having a dynamic game does not necessarily lead to the selection of MPE. Most subjects who do not use Markov strategies aim to 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.

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

[42]Our experiments do not allow us to identify whether this comes from the presence of the very high state (for instance, causing cooperation to unravel from the top) or the very-low state (where lower possible payoffs cause players to focus on the individually rational actions).

[43]In the appendix we provide estimates from the SFEM procedure only over the four Markov strategy profiles, using data from the last three periods of each supergame. While the pivot estimate has approximately half of the data on M_{DD} and M_{CD}, the En-DPD-$\tilde{\Theta}$ has 100 percent across the two.

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, En-DCP-E and En-DPD-$\tilde{\Theta}$. 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 more-lenient history-dependent strategies.[44] 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 outcome increases the selection of more-efficient cooperative outcomes.*

Evidence: In En-DPD-CC and En-DPD-HT we make it easier to sustain joint-cooperation, reducing the temptation to defect from an S_{CD} trigger. In both cases, cooperation increases, though more so for the dynamic modification in En-DPD-HT.

Result 4 (Response to Static/Dynamics)**.** *Behavior is sensitive to both static and dynamic tensions within the game.*

Evidence: Theory motivates that both static and dynamic effects will drive whether an outcome is an equilibrium. Static changes to the stage-game payoff $u_i(a, \theta)$ produce significant shifts in behavior (En-DPD$\rightarrow${En-DPD-CC, En-DCP-E, En-DCP-M, En-DPD-X}). Similarly, changes to the game's transition rule affect the game's continuation value, and we again see significant shifts in behavior (En-DPD$\rightarrow${En-DPD-HT,Ex-DPD, Ex-SPD,En-DPD-$\tilde{\Theta}$}).

Result 5 (Complexity)**.** *Adding exogenous states (shocks) does not lead to an increase in MPE play, and more cooperative outcomes are still common. However, while adding more endogenous states does not lead to an increase in the selection of the MPE, the presence of additional states does alter behavior, leading to lower cooperation and greater miscoordination.*

Evidence: Our two treatments with richer state-spaces lead to differing rates of cooperation. Where we add exogenous non-persistent shocks to the payoffs each round (En-DPD$\rightarrow$En-DPD-X)

[44] Along the equilibrium path strategies that implement joint cooperation and M_{CC} are identical and the SFEM cannot separately identify them. An simple alternative to evaluate whether what we identify as M_{CC} is actually the successful outcome of joint cooperation is to look at subjects' behavior across supergames. We first identify all subjects who have at least one supergame where the subject and their partner cooperated in every period. For these "cooperating" subjects we then focus on all other supergames where they started by cooperating but their partner defected at least once. In these supergames behavior would be consistent with M_{CC} if subjects keep on cooperating regardless of their partner's behavior. For the treatments in which M_{CC} is statistically significant in Table 5 we report the proportion of supergames where the choices of "cooperating" subjects can still be captured by M_{CC} even if their partner defected at least once. The proportions are: 25.7, 29.2, 18.6 and 23.1 in En-DPD-CC, En-DPD-HT, En-SPD ($\theta = L$) and En-DPD-X, respectively. This figures indicate that the vast majority of subjects who use M_{CC} once do punish if their partner defects.

the aggregate observed behavior looks similar, if anything moving away from the MPE and towards higher-efficiency outcomes. When we add additional endogenous states (En-DPD$\rightarrow$En-DPD-Θ) overall Markov play stays fairly constant, but the lower-payoff Markov profile M_{DD} is selected at a relatively higher frequency, and M_{CC} is much reduced. As the number of states increase, so too does the set of plausible strategies. Coordination therefore becomes more challenging, and some subjects become more pessimistic, pushing behavior away from cooperation.

7.2. **Toward a Selection Index.** The larger experimental literature on infinitely repeated games has identified 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 the set of beliefs on the other player being a conditional cooperator that would 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. As a simple rule of thumb, the literature offers the binary selection criterion: if the grim-trigger in the particular environment is risk-dominant (basin larger than one half), history-dependent cooperative SPE are more likely.

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 extension of the basin over? 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 results in the last subsection we motivate the following: i) The basin calculation should respond to *both* static and dynamic strategic externalities, motivating extensions of the BA that integrate the entire dynamic game machinery into their calculation; ii) symmetric strategies are focal; iii) where the MPE is selected, the best-case 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.

The above motivates our focus on a binary strategy selection across: i) the dynamic-game Γ's most-efficient symmetric MPE (M_Γ), and ii) the most-efficient symmetric outcome path sustainable as an SPE with a reversion to M_Γ on any deviation (S_Γ). 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

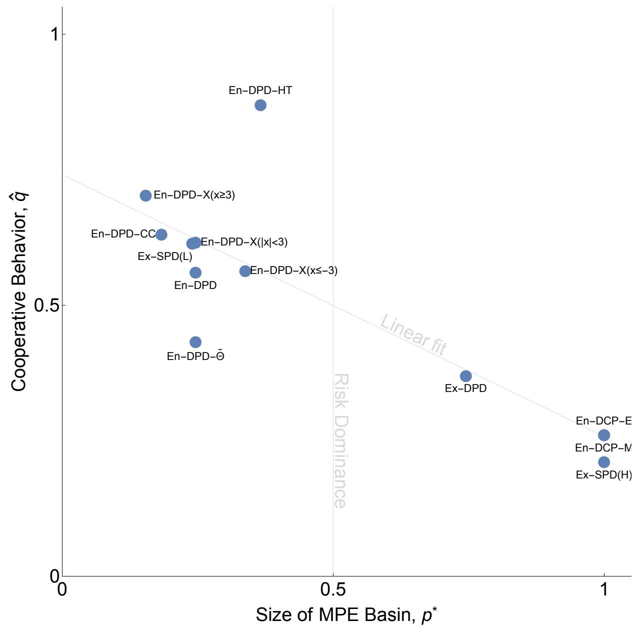

FIGURE 2. Basin of Attraction for the MPE

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, $p^\star(S_\Gamma, M_\Gamma; \Gamma)$. The vertical axis 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).

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.[45]

Given the theoretical index $p^\star_\Gamma$ we now want to compare the index with a measure of behavior in the experimental sessions, $\hat{q}_\Gamma$. In the infinitely-repeated game literature the focal outcome measure is the first-period cooperation rate in supergames. Again we need to make a modification to account for the dynamic-game setting: In the pivot, both the MPE and the SPE strategies are predicted to cooperate in the first low period. So first-round cooperation will not be informative on the differential selection between the MPE and more-cooperative outcomes. Instead we focus on the cooperation rate $\hat{q}_\Gamma$ in the first round where M_Γ and S_Γ are expected to choose differing actions. For the pivot, the strategies S_{CD} and M_{CD} cooperate in round one and then choose differing actions in the high state for the second-round, C and D, respectively. Looking at the last five supergames in the pivot

[45]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.

sessions, 56.2 percent of our subject-supergames have paths consistent with $((C,C),L),(C,\cdot,H)$ and so for the pivot the basin-behavior pair $(p^{\star},\hat{q})$ is given by $(0.246,0.562)$.[46] Figure 2 provides a plot of each index-behavior point across our treatment set (and some sub-treatments where the prediction differs) illustrating the basin's predictive power.[47]

As illustrated by Figure 2, the basin calculations predict initial supra-MPE cooperation fairly well. Given the 12 data points represented, an OLS regression indicates that the cooperation rates are significantly related to the size of the MPE basin for each game (99 percent confidence). For a simpler rule, risk-dominance of the MPE is also predictive of selection. A simple, easy-to-calculate criterion for the appropriateness of an MPE assumption is whether the MPE risk dominates the best-case symmetric SPE (with the MPE triggered as a history-dependent punishment on any deviation).

However, rather than focus here on the overall predictive success of the basin, we now instead outline where it might be further refined through more-targeted future research. The largest disconnect between the fitted relationship for the MPE basin measure and cooperation are in En-DPD-HT and En-DPD-$\tilde{\Theta}$. Both treatments provide insights into ways in which the very simple basin calculation defined here might be modified.

For En-DPD-HT the index predicts *decreased* cooperation relative to the pivot, as the basin calculation is $p^{\star}(S_{CD},M_{CD})=0.366$. In contrast, this is the treatment with the most cooperative outcomes, where 87.1 percent of supergames have high-state cooperation in period two. One reason for this treatment being an outlier in the figure is that subjects in this treatment coordinate on much-harsher punishments than the basin calculation allows for. Given a deviation, subjects in En-DPD-HT respond with a switch to the worst-case MPE M_{DD} (rather than the modeled M_{CD} response).

---

[46]In this way, we are more conservative in ascribing behavior as cooperative, as a player may have been cooperative in period one, but was defected on. Our behavior measure $\hat{q}_{\Gamma}$ only incorporates cooperation with a consistent path up to the predicted point.

[47]While in some treatments the same basin calculation between S_{CD} and M_{CD} is used (in particular En-DPD-CC, En-DPD-HT and En-DPD-$\tilde{\Theta}$) 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 2 we aggregate the shocks into three categories ($x_2\leq -3$, $-3<x_2<3$, and $3\leq x_2$), and plots the basin-behavrior pairs. For Ex-DPD the basin calculation shifts to account for changes in both the MPE and best-case symmetric SPE. The MPE prediction in this game shifts to M_{DD} because of the exogenous transitions. Additionally, S_{DD} is not an SPE here, and our basin calculation calls for the best symmetric SPE. Instead we use the symmetric trigger that supports cooperation in low and defection in high with an M_{DD} trigger on any deviation (call this strategy X_{DD}). In the Ex-SPD treatment where the low state is initially selected the basin calculation is the standard infinitely-repeated PD game calculation $p^{\star}(S_D,M_D)$. Finally, in three treatments the best-*symmetric* SPE *is* the best-case MPE, in which case the basin for the MPE is the full set of beliefs, with measure one. 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, so Figure 2 reflects this with $\hat{q}_{\Gamma}$ reflecting the period-one cooperation rate.

The basin calculation for cooperation relative to this worst-case MPE is $p^{\star}(S_{DD}, M_{DD}) = 0.071$, and so the treatment would be much less of an outlier. A desirable modification to our simple basin calculation might initially compare coordination across S_{CD}, S_{DD}, M_{CD} and M_{DD} and discern that the $p^{\star}(S_{DD}, M_{DD})$ comparison was the most relevant margin, thereby eliminating the best-case MPE M_{CD}.[48]

Our second outlier is En-DPD-$\tilde{\Theta}$, where the basin is $p^{\star}(S_{CD}, M_{CD}) = 0.246$, which is identical to the pivot's basin by construction. But, as detailed in Section 6, the additional endogenously reachable states lead to greater subject pessimism at the very start of the game, and less coordination on cooperative strategies. This can be seen in Figure 2 with the fall in high-state cooperation in round two from 56 percent in the pivot to 43.3 percent in En-DPD-$\tilde{\Theta}$. This suggests that as the number of endogenous states increases coordination on cooperation is more demanding, which is an aspect that is not currently captured by our index.[49] Future research may help pin down more-general indices that incorporate the effects of richer state complexity on coordination.

8. Conclusion

Our paper explores a set of nine dynamic games under an infinite-time horizon. While many applications of dynamic games focus on Markov-perfect equilibria, our results suggest that the selection of state-dependent strategies depends on features of the game. Our core treatment is a simple two-state extension of the infinitely repeated prisoner's dilemma, and we find behavior that is conceptually closer to the experimental literature on repeated games than the theoretically focal MPE assumption. Most behavior is consistent with history-dependent strategies that aim to achieve greater efficiency than the MPE prediction.

Our treatments also allow us to identify conditions under which Markov play may become more prominent. First, we find that the MPE prediction is more frequent in games where coordination on history-dependent strategies is harder. This happens as we weaken the strategic externalities through changes to the dynamics and stage game payoffs, and as we decrease the focality of symmetric history-dependent outcomes. Second, we do find that changes to the complexity of the state-space can lead to breakdowns of cooperation. Though small with respect to the frequency strategies are initially selected, in the long-run the effect is more pronounced, as many more supergames are miscoordinated through greater variation in strategic choices. This finding might help

[48] One reason for a change in focus might also be changes to the individually rational action in the high state for En-DPD-HT.

[49] Where the basin calculation in En-DPD-$\tilde{\Theta}$ has all strategic uncertainty realized in period two, there are higher-powered coordination challenges in the very high state. To see this, if we started the En-DPD-$\tilde{\Theta}$ game in the very-high state, we would get the much-larger MPE basin $p^{\star}(S_{CD}, M_{CD}; \theta_1 = vH) = 0.889$. Richer dynamic games may have ongoing strategic uncertainty if the game is constantly entering new states, which may lead to an unraveling towards the MPE.

explain the greater selection of symmetric MPE in other large state-space experiments on dynamic games (Battaglini et al., 2012, 2014; Vespa, 2015), where increased strategic uncertainty from the many possible future states pushes the game towards state-dependent MPE behavior. However, when the state-space is increased through exogenous, non-persistent shocks (which are common in many industrial organization applications) we find a small increase in history dependence. This is also is consistent with other experimental games with a larger state-space (see Salz and Vespa, 2015).

While our results allow us to bridge earlier findings in repeated and dynamic games, our relatively large number of treatments illustrate a richness in subject behavior. That more-efficient history-dependent strategies emerge in our lab data suggests researchers should be somewhat wary of making Markov assumptions. If incentive-compatible strategies with Pareto superior outcomes are quickly learned and deployed by undergraduate students matched anonymously with one another in the lab, it is hard to believe they are will not be present in the field, where participants engage in longer interactions and with more channels for coordination. Future research can further explore and pin down what drives selection, while many other 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.

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.

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.

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.

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.

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.

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," 2015. 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 the Dynamic Common Pool Game," 2015. UCSB working paper.

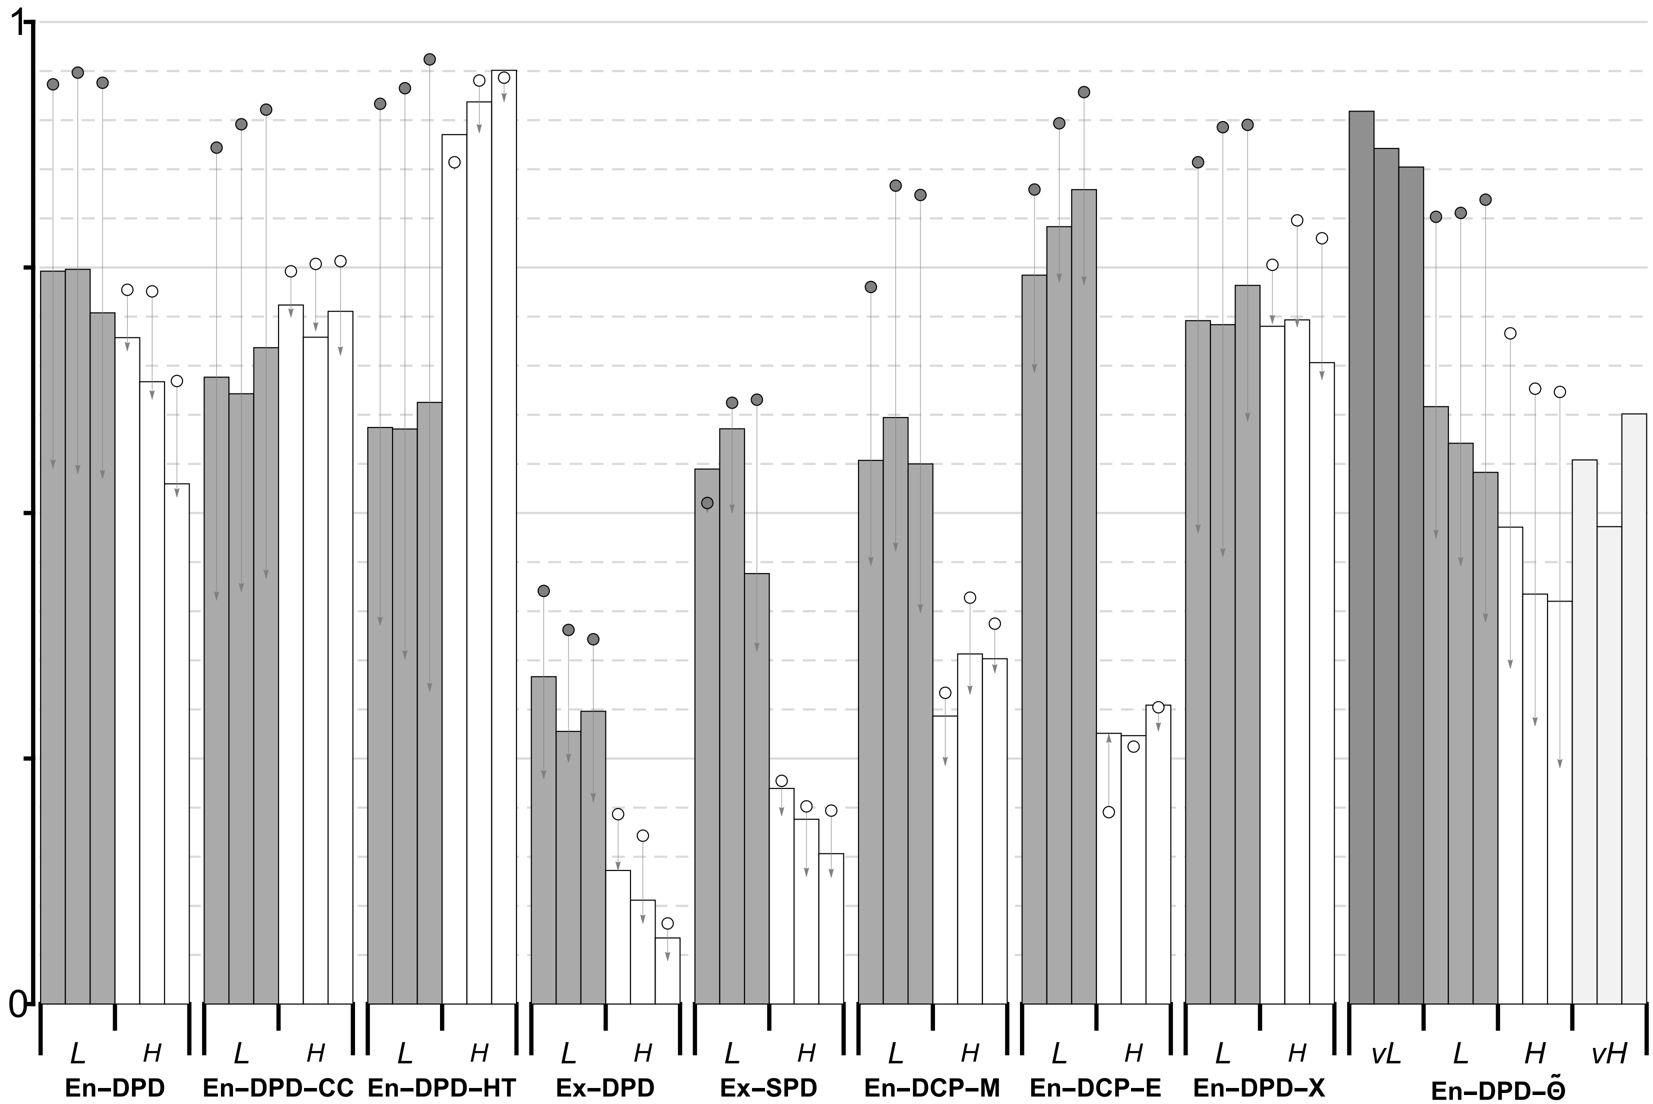

FIGURE 3. Cooperation Rates, by treatment/state/supergame-block

Note: 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.

TABLE 5. Strategy Frequency Estimation Method Output: Last Five Supergames

Strategies	En-DPD	En-DPD-CC	En-DPD-HT	Ex-DPD	Ex-SPD ($\theta = L$)	Ex-SPD ($\theta = H$)	En-DCP-M	En-DCP-E	En-DPD-X	En-DPD-$\tilde{\Theta}$
Markov										
M_{CC}	0.117	0.173**	0.251**	0.000	0.112*	0.009	0.068	0.092	0.347***	0.023
	(0.072)	(0.076)	(0.116)	(0.013)	(0.060)	(0.014)	(0.053)	(0.059)	(0.091)	(0.044)
M_{DD}	0.024	0.039	0.023	0.582***	0.246***	0.523***	0.077	0.048	0.027	0.166*
	(0.030)	(0.039)	(0.017)	(0.145)	(0.082)	(0.098)	(0.061)	(0.038)	(0.034)	(0.087)
M_{CD}	0.212*	0.057	0.041	0.073*			0.279*	0.651***	0.138*	0.167*
	(0.127)	(0.063)	(0.033)	(0.041)			(0.158)	(0.181)	(0.081)	(0.097)
M_{DC}	0.000	0.000	0.000	0.000			0.063	0.000	0.000	0.000
	(0.003)	(0.010)	(0.026)	(0.000)			(0.043)	(0.002)	(0.015)	(0.038)
History-dependent										
S_{DD}	0.106	0.227*	0.479***	0.180*	0.184*	0.078	0.039	0.000	0.069	0.265***
	(0.095)	(0.119)	(0.133)	(0.109)	(0.102)	(0.061)	(0.048)	(0.009)	(0.066)	(0.076)
S_{CD}	0.206**	0.075	0.000	0.045			0.070	0.088*	0.245**	0.106
	(0.085)	(0.067)	(0.043)	(0.033)			(0.057)	(0.052)	(0.110)	(0.065)
TfT	0.254***	0.304***	0.182	0.032	0.324***	0.131*	0.139*	0.069	0.089	0.129
	(0.082)	(0.093)	(0.125)	(0.052)	(0.080)	(0.072)	(0.080)	(0.078)	(0.063)	(0.088)
$STfT$	0.023	0.016	0.000	0.065	0.134	0.258	0.000	0.000	0.021	0.000
	(0.027)	(0.020)	(0.002)	(0.062)			(0.008)	(0.002)	(0.031)	(0.003)
A_{DD}	0.059	0.046	0.024	0.022			0.164	0.051	0.029	0.101
	(0.039)	(0.041)	(0.029)	(0.043)			(0.111)	(0.050)	(0.034)	(0.067)
A_{CD}	0.000	0.065	0.000	0.000			0.102	0.000	0.035	0.043
γ	0.643***	0.529***	0.364***	0.347***	0.532***	0.451***	0.706***	0.526***	0.635***	0.669***
	(0.058)	(0.071)	(0.041)	(0.036)	(0.048)	(0.051)	(0.063)	(0.080)	(0.089)	(0.062)
β	0.826	0.869	0.940	0.947	0.868	0.902	0.805	0.870	0.828	0.817

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

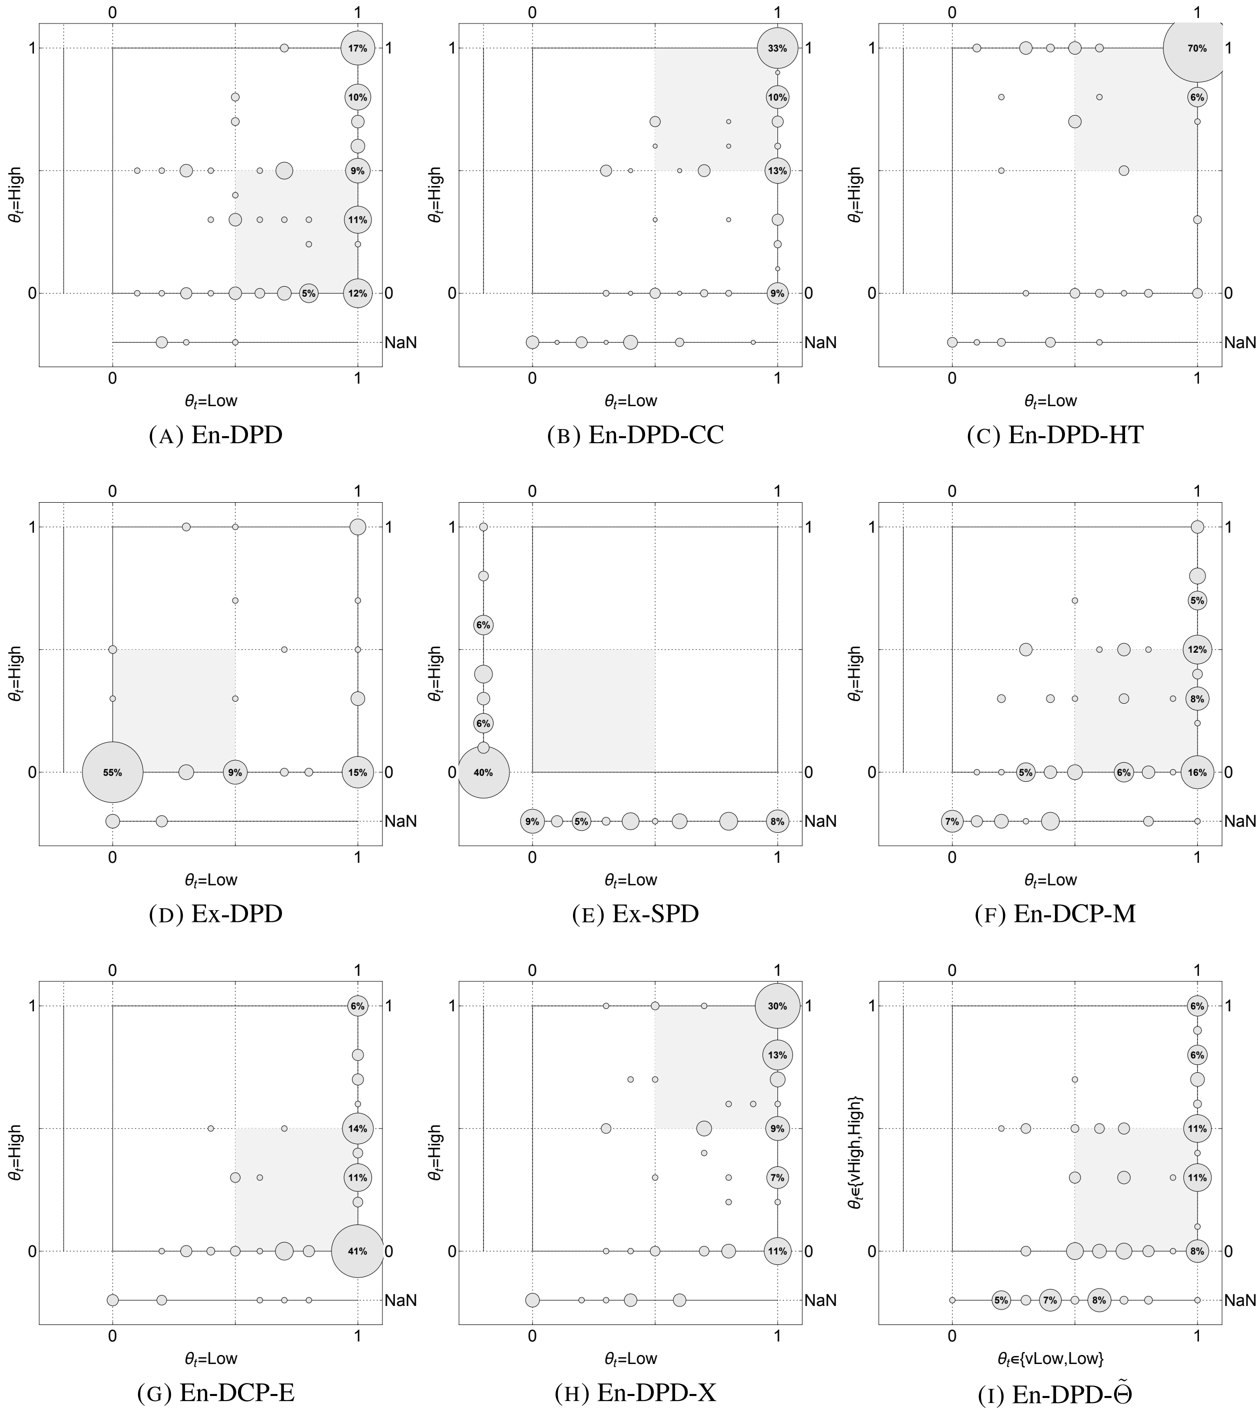

(A) En-DPD (B) En-DPD-CC (C) En-DPD-HT

(D) Ex-DPD (E) Ex-SPD (F) En-DCP-M

(G) En-DCP-E (H) En-DPD-X (I) En-DPD-$\tilde{\Theta}$

FIGURE 4. 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.

APPENDIX A. SUPPLEMENTARY MATERIAL: FIGURES AND TABLES

Tables 6 and 7 present the stage games for the En-DPD-CC and En-DPD-X treatments, respectively.

TABLE 6. En-DPD-CC Stage Games

θ=Low

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

θ=High

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

TABLE 7. En-DPD-X Stage games

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

1:	2: C	2: D
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 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)	$(\star\star\star)$
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)	$(\star\star\star)$
En-DPD-X	0.856	(0.036)	$(\star)$	0.635	(0.055)	$(\star\star\star)$
En-DPD-$\tilde{\Theta}$	0.453	(0.033)	$(\star\star\star)$	0.254	(0.032)	$(\star\star\star)$

Note: 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 3 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 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.

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

	En-DPD-CC	En-DPD-HT	Ex-DPD	Ex-SPD	En-DCP-M	En-DCP-E	En-DPD-X	En-DPD-$\tilde{\Theta}$
Constant	0.956***	0.956***	0.987***	0.971***	0.984***	0.960***	0.963***	0.942***
	(0.028)	(0.027)	(0.035)	(0.026)	(0.023)	(0.026)	(0.025)	(0.031)
Treatment	-0.027	0.023	-0.566***	-0.331***	-0.114***	-0.009	-0.042	-0.119*
	(0.048)	(0.036)	(0.112)	(0.106)	(0.032)	(0.040)	(0.050)	(0.064)
State	-0.331***	-0.333***	-0.335***	-0.336***	-0.329***	-0.332***	-0.330***	-0.331
	(0.057)	(0.057)	(0.057)	(0.056)	(0.059)	(0.057)	(0.058)	(0.057)
State × Treatment	0.136*	0.294***	0.067	-0.052	-0.169***	-0.318***	0.163***	0.040
	(0.075)	(0.067)	(0.089)	(0.061)	(0.065)	(0.089)	(0.059)	(0.070)
Supergame 12	-0.017	-0.011	-0.015	-0.017	-0.027	-0.031**	-0.010	0.003
	(0.031)	(0.023)	(0.038)	(0.031)	(0.026)	(0.015)	(0.022)	(0.027)
Supergame 13	-0.024*	-0.029*	-0.039	-0.069***	-0.074***	-0.011	-0.035	0.019
	(0.013)	(0.015)	(0.033)	(0.026)	(0.027)	(0.023)	(0.025)	(0.046)
Supergame 14	-0.017	-0.026	-0.087***	-0.042	-0.076***	-0.042*	-0.032	-0.030
	(0.039)	(0.029)	(0.032)	(0.028)	(0.020)	(0.022)	(0.036)	(0.027)
Supergame 15	-0.034	-0.024	-0.104*	-0.037	-0.055***	-0.026	-0.047***	-0.017
	(0.021)	(0.021)	(0.062)	(0.032)	(0.017)	(0.025)	(0.016)	(0.025)

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

Note: 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 × 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	θ =**Low**			θ =**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)	
Ex-DPD	0.124	(0.050)	$(\star\star\star)$	0.049	(0.022)	$(\star\star)$
Ex-SPD	0.286	(0.068)	$(\star\star\star)$	$0.040^{\dagger}$	(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)	
En-DPD-X	0.256	(0.065)	$(\star\star\star)$	0.256	(0.055)	$(\star\star\star)$
En-DPD-$\tilde{\Theta}$	0.421	(0.041)	$(\star\star\star)$	0.219	(0.030)	$(\star\star\star)$

Note: 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. SFEM Output: Including only the last three periods of each history (Last 5 supergames)

Strategies	**En-DPD**	**En-DPD-CC**	**En-DPD-HT**	**Ex-DPD**	**En-DCP-M**	**En-DCP-E**	**En-DPD-X**	**En-DPD-$\tilde{\Theta}$**
M_{CC}	0.194	0.175	0.318***	0.045	0.015	0.126*	0.391***	0.001
	(0.136)	(0.162)	(0.118)	(0.040)	(0.062)	(0.067)	(0.129)	(0.055)
M_{DD}	0.165	0.136	0.000	0.874***	0.407***	0.103	0.041	0.603***
	(0.110)	(0.085)	(0.024)	(0.070)	(0.144)	(0.077)	(0.056)	(0.135)
M_{CD}	0.380***	0.135*	0.048	0.081	0.432***	0.771***	0.301***	0.396***
	(0.132)	(0.072)	(0.030)	(0.054)	(0.110)	(0.110)	(0.089)	(0.124)
M_{DC}	0.262	0.555	0.634	0.000	0.146	0.000	0.267	0.000
γ	1.065***	0.806***	0.408***	0.362***	0.888***	0.615***	0.819***	1.131***
	(0.119)	(0.085)	(0.070)	(0.060)	(0.098)	(0.133)	(0.099)	(0.159)
β	0.719	0.776	0.921	0.941	0.755	0.836	0.772	0.708

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

TABLE 12. Path for the State: Last Five Supergames

Sequence	En-DPD	En-DPD-CC	En-DPD-HT	Ex-DPD	Ex-SPD	En-DCP-M	En-DCP-E	En-DPD-X	En-DPD-$\tilde{\Theta}$
LLLLL	0.038	0.121	0.057	0.048	0.371	0.200	0.048	0.086	0.257
	(0.004)	(0.009)	(0.005)	(0.004)	(0.023)	(0.016)	(0.004)	(0.008)	(0.019)
LHLHL	0.038	0.014	0.029	0.00		0.114	0.362	0.019	0.067
	(0.004)	(0.001)	(0.003)	–		(0.010)	(0.023)	(0.002)	(0.006)
LHHHH	0.476	0.607	0.733	0.371		0.19	0.267	0.600	0.314
	(0.024)	(0.020)	(0.019)	(0.023)		(0.015)	(0.019)	(0.023)	(0.021)
LHHLL	0.124	0.086	0.010	0.029		0.067	0.01	0.048	0.105
	(0.011)	(0.007)	(0.001)	(0.003)		(0.006)	(0.001)	(0.004)	(0.009)
LHHLH	0.105	0.036	0.019	0.038		0.095	0.143	0.067	0.057
	(0.009)	(0.003)	(0.002)	(0.004)		(0.008)	(0.012)	(0.006)	(0.005)
LHHHL	0.086	0.064	0.057	0.143		0.076	0.048	0.057	0.095
	(0.008)	(0.005)	(0.005)	(0.012)		(0.007)	(0.004)	(0.005)	(0.008)
LHLLL	0.048	0.021	0.067	0.000		0.057	0.010	0.000	0.029
	(0.004)	(0.002)	(0.006)	–		(0.005)	(0.001)	–	(0.003)
HHHHH					0.629				
					(0.023)				
All Other	0.086	0.05	0.029	0.371	0.000	0.200	0.114	0.124	0.076
	(0.008)	(0.004)	(0.003)	(0.023)		(0.016)	(0.010)	(0.011)	(0.007)

Note: Data for En-DPD-$\tilde{\Theta}$ merges vL and L, and vH and H.

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

Treatment	5 or more observed Supergames				
En-DPD	CC,**DC**,**DD**,CC,**DD**	CC,**CC**,**CC**,**CC**,**CC**	CC,**CC**,**CC**,**CC**,**DC**	CC,**DC**,**DC**,**DC**,**DC**	
	0.095	0.086	0.057	0.057	
En-DPD-CC	CC,**CC**,**CC**,**CC**,**CC**	CC,**CC**,**CC**,**DC**,**DD**	CC,**CC**,**CC**,**CC**,**DC**	CC,**DC**,**DD**,CD,CD	
	0.264	0.079	0.050	0.043	
En-DPD-HT	CC,**CC**,**CC**,**CC**,**CC**	CC,**CC**,**CC**,**CC**,**DC**			
	0.610	0.124			
Ex-DPD	DD,DD,DD,DD,DD	DC,DD,DD,DD,DD			
	0.381	0.352			
Ex-SPD (Low)	CC,CC,CC,CC,DC	DC,DD,DD,DD,DD	CC,CC,CC,CC,CC		
	0.128	0.128	0.128		0.629
Ex-SPD (High)	**DD**,**DD**,**DD**,**DD**,**DD**	**DC**,**DC**,**DD**,**DD**,**DD**	**DC**,**DD**,**DD**,**DD**,**DD**		
	0.515	0.091	0.076		
En-DCP-M	CC,**DC**,**CD**,**DC**,**CD**	CC,**DC**,**DD**,CC,**DD**	CC,**DD**,CC,**DD**,CC	CC,**DD**,CC,**DD**,DC	DC,DC,DD,DD,DD
	0.076	0.067	0.057	0.057	0.057
En-DCP-E	CC,**DD**,CC,**DD**,CC	CC,**DC**,**DD**,CC,**DD**	CC,**DC**,**CD**,**DC**,**CD**	CC,**DD**,CC,**DD**,DC	CC,**DC**,**DC**,**DC**,**DC**
	0.257	0.105	0.095	0.095	0.048
En-DPD-X	CC,**CC**,**CC**,**CC**,**CC**	CC,**CC**,**CC**,**CC**,**DC**	CC,**DC**,**DC**,**DC**,**DC**		
	0.162	0.095	0.048		
En-DPD-$\tilde{\Theta}$	CC,**DD**,CC,**DD**,DC	DC,DD,CC,DD,CC	CC,**DC**,**DD**,CC,**DD**	CC,**CC**,CC,CC,CC	
	0.057	0.057	0.048	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 5 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 19). 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 5 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. A similar effect (albeit smaller) is present for En-DPD-$\tilde{\Theta}$. Other than these discrepancies the estimates reported in Table 5 remain largely unchanged.

Table 16 presents estimates when we expand the set of Markov strategies in treatments where we change the size of the state-space. To explain the extra strategies included for En-DPD-X, consider first Figure 5. 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 5. 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 5, 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.

For En-DPD-$\tilde{\Theta}$, we explored sequences of actions and states (see Tables 12 and 13) that can be rationalized with Markov strategies that do not prescribe the same choice for vH and H, and vL and L. We include two additional strategies. M_{CDDD}, which cooperates only in vL and defects otherwise and M_{CCCD} that only defects in vH. The estimates in Table 16 show that both strategies capture approximately 10 percent of the mass and are statistically significant. The goodness-of-fit measure relative to Table 5 increases by 6 points. The strategy that suffers the largest loss of mass (relative to Table 5) is M_{DD}, which is now at zero.

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

Strategies	En-DPD	En-DPD-CC	En-DPD-HT	Ex-DPD	Ex-SPD ($\theta = L$)	Ex-SPD ($\theta = H$)	En-DCP-M	En-DCP-E	En-DPD-X	En-DPD-$\tilde{\Theta}$
Markov										
M_{CC}	0.118*	0.177**	0.206*	0.000	0.104*	0.000	0.066	0.092	0.362***	0.016
	(0.072)	(0.080)	(0.122)	(0.015)	(0.060)	(0.014)	(0.053)	(0.059)	(0.092)	(0.045)
M_{DD}	0.045	0.087	0.032	0.630***	0.268***	0.733***	0.265**	0.069	0.027	0.238**
	(0.042)	(0.058)	(0.041)	(0.137)	(0.092)	(0.118)	(0.122)	(0.054)	(0.041)	(0.096)
M_{CD}	0.265	0.115	0.050	0.093**			0.393	0.682	0.189	0.275***
	(0.102)	(0.060)	(0.033)	(0.047)			(0.092)	(0.109)	(0.071)	(0.087)
History-dependent										
S_{DD}				0.194	0.168	0.080				
				(0.109)	(0.100)	(0.061)				
S_{CD}	0.207**	0.194**	0.041				0.077	0.088*	0.301**	0.304***
	(0.082)	(0.080)	(0.078)				(0.062)	(0.052)	(0.122)	(0.087)
TfT	0.365	0.427	0.671	0.083	0.460	0.187	0.198	0.069	0.121	0.168
γ	0.679***	0.571***	0.410***	0.357***	0.564***	0.483***	0.805***	0.547***	0.664***	0.775***
	(0.061)	(0.085)	(0.056)	(0.042)	(0.054)	(0.051)	(0.094)	(0.089)	(0.093)	(0.071)
β	0.813	0.852	0.920	0.943	0.855	0.888	0.776	0.862	0.819	0.784

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-DCP-M**	**En-DCP-E**	**En-DPD-X**	**En-DPD-$\tilde{\Theta}$**
			Markov			
M_{CC}	0.117*	0.173**	0.068	0.092	0.347***	0.000
	(0.071)	(0.077)	(0.052)	(0.058)	(0.090)	(0.040)
M_{DD}	0.024	0.039	0.077	0.048	0.027	0.167*
	(0.030)	(0.039)	(0.061)	(0.038)	(0.034)	(0.087)
M_{CD}	0.211*	0.057	0.311**	0.651***	0.138*	0.194**
	(0.127)	(0.063)	(0.157)	(0.180)	(0.081)	(0.090)
M_{DC}	0.000	0.000	0.063	0.000	0.000	0.000
	(0.003)	(0.010)	(0.042)	(0.001)	(0.014)	(0.038)
			History-dependent			
S_{DD}	0.106	0.227*	0.000	0.000	0.069	0.239***
	(0.102)	(0.124)	(0.042)	(0.008)	(0.065)	(0.074)
S_{CD}	0.000	0.075	0.000	0.088	0.245**	0.045
	(0.076)	(0.069)	(0.052)	(0.056)	(0.123)	(0.054)
TfT	0.232***	0.304***	0.134*	0.069	0.089	0.068
	(0.083)	(0.094)	(0.079)	(0.078)	(0.061)	(0.079)
$STfT$	0.023	0.016	0.000	0.000	0.021	0.000
	(0.027)	(0.020)	(0.009)	(0.002)	(0.030)	(0.003)
A^D_{DD}	0.059	0.046	0.161	0.051	0.029	0.099
	(0.039)	(0.041)	(0.109)	(0.050)	(0.035)	(0.065)
A^D_{CD}	0.000	0.065	0.080	0.000	0.035	0.000
	(0.103)	(0.060)	(0.130)	(0.143)	(0.054)	(0.039)
A^C_{DD}	0.000	0.000	0.044	0.000	0.000	0.063
	(0.058)	(0.076)	(0.038)	(0.002)	(0.055)	(0.057)
A^C_{CD}	0.228	0.000	0.061	0.000	0.000	0.126
γ	0.641***	0.529***	0.701***	0.526***	0.635***	0.634***
	(0.058)	(0.071)	(0.064)	(0.080)	(0.088)	(0.044)
β	0.826	0.869	0.806	0.870	0.828	0.829

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

TABLE 16. SFEM Output: Additional Strategies in Complexity Treatments

Strategies	En-DPD-X	En-DPD-$\tilde{\Theta}$
	Markov	
M_{CC} (M_{CCCD})	0.253***	0.000
	(0.078)	(0.041)
M_{DD} (M_{CDDD})	0.027	0.166*
	(0.034)	(0.088)
M_{CD} (M_{CCDD})	0.133*	0.176*
	(0.071)	(0.092)
M_{DC} (M_{CDCD})	0.000	0.000
	(0.013)	(0.039)
M_{CCCC}		0.059
		(0.077)
M_{DDDD}		0.000
		(0.038)
$M^{x}_{CCC,DCC}$	0.203**	
	(0.098)	
$M^{x}_{CCC,DDC}$	0.002	
	(0.048)	
	History-dependent	
S_{DD} (S_{CDDD})	0.073	0.266***
	(0.062)	(0.075)
S_{CD} (S_{CCDD})	0.162	0.109*
	(0.119)	(0.065)
$S^{x}_{CCC,DCC}$	0.000	
	(0.019)	
$S^{x}_{CCC,DDC}$	0.000	
	(0.020)	
TfT	0.063	0.091*
	(0.056)	(0.055)
$sTfT$	0.015	0.000
	(0.024)	(0.003)
A_{DD} (A_{CDDD})	0.032	0.102
	(0.036)	(0.066)
A_{CD} (A_{CCDD})	0.038	0.031
γ	0.588***	0.645***
	(0.070)	(0.054)
β	0.846	0.825

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

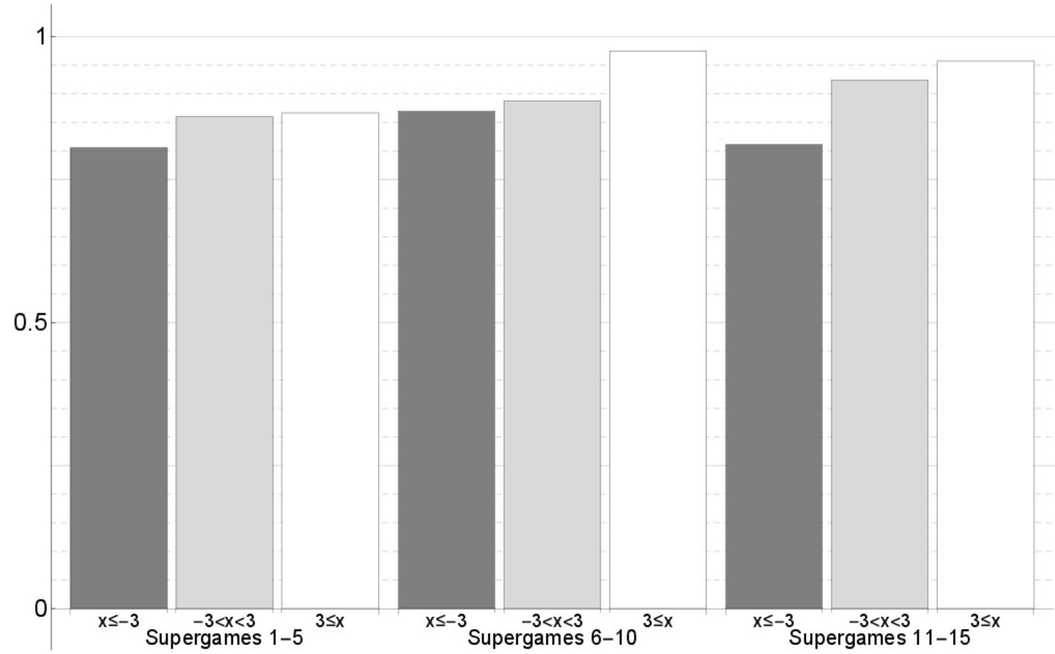

(A) Low state, period one

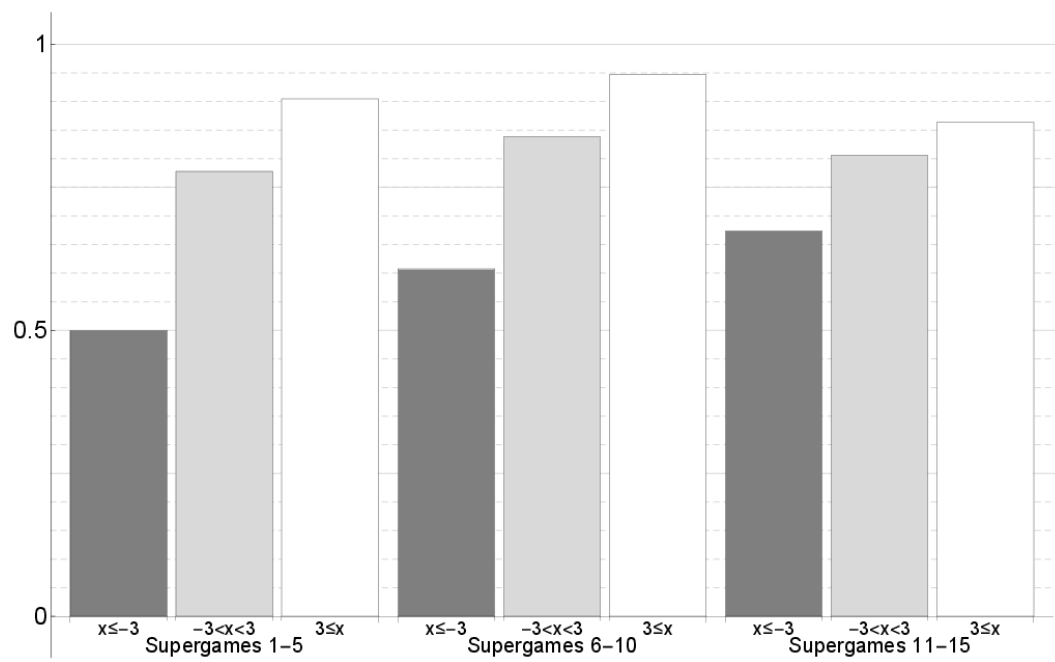

(B) High state, period two

FIGURE 5. Cooperation rates in En-DPD-X

Note: 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).