

Vespa, Emanuel; Wilson, Alistair J.

Working Paper

Dynamic Incentives and Markov Perfection: Putting the 'Conditional' in Conditional Cooperation

CESifo Working Paper, No. 5413

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Vespa, Emanuel; Wilson, Alistair J. (2015) : Dynamic Incentives and Markov Perfection: Putting the 'Conditional' in Conditional Cooperation, CESifo Working Paper, No. 5413, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/113746>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Working Papers

www.cesifo.org/wp

Dynamic Incentives and Markov Perfection: Putting the 'Conditional' in Conditional Cooperation

Emanuel Vespa
Alistair J. Wilson

CESIFO WORKING PAPER NO. 5413

CATEGORY 13: BEHAVIOURAL ECONOMICS

JUNE 2015

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: www.CESifo-group.org/wp

ISSN 2364-1428

CESifo

Center for Economic Studies & Ifo Institute

Dynamic Incentives and Markov Perfection: Putting the ‘Conditional’ in Conditional Cooperation

Abstract

This paper experimentally examines the selection of equilibria in dynamic games. Our baseline treatment is a two-state extension of an indefinitely repeated prisoner’s dilemma, which we modify in series of treatments to study the focality of efficiency and symmetry, the effect dynamic and static strategic externalities, and the size of the state-space. Subjects in our experiments show an affinity for conditional cooperation, readily conditioning their behavior on both the state of the world, and recent history of play. With strong dynamic and static externalities present we see most subjects coordinate on efficiency by conditioning on past play. However, when we remove either type of strategic externality, conditioning on just the state becomes more common, and behavior is consistent with the Markov-perfect prediction. Changes to the environment’s state-space are more nuanced: perturbations of the game with small-sized noise does not lead to more state-conditioned behavior; however, a richer set of endogenous states does lead to more Markov-perfect behavior.

JEL-Code: C730, C920, D030, D900.

Keywords: dynamic cooperation, equilibrium selection, history dependence.

Emanuel Vespa
University of California
Department of Economics
2127 North Hall
USA - Santa Barbara, CA 93106
vespa@ucsb.edu

Alistair J. Wilson
Department of Economics
University of Pittsburgh
4907 Wesley W. Posvar Hall
230 South Bouquet Street
USA - 15260 Pittsburgh PA
alistair@pitt.edu

Date: June, 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, and the 2015 Conference on Social Dilemmas.

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 put enough weight on long-run benefits from sustained cooperation, threats to condition future behavior on present outcomes and thereby punish opportunistic choices can be credible. 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 observed 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.

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

¹A 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).

²Here 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.

games ignores such conditioning. Restricting attention to Markov strategies can rule out efficient outcomes, even where they are supportable in an SPE.

The available experimental evidence on human 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 study the connection between the experimental literatures on infinitely repeated and dynamic games. Given the two distinct sets of results, a natural task is characterizing which properties of a dynamic game might lead to the selection of state-dependent strategies, and which to less-restrictive history-dependent ones. 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 miscoordination 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. The game we use 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 the 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.

³Battaglini 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 (2014), 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 (2014).

In the pivot (and many of our modified versions of it) we find that a majority of human 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 change the efficient frontier in the dynamic game, making a single symmetric SPE more focal while holding constant the MPE prediction. These treatments allow us to study whether behavior responds in the direction of greater use of history-dependent strategies. While our pivot has many SPE outcomes with higher payoffs than the symmetric MPE, our efficiency manipulations will make one of these outcomes (joint cooperation in both states) more focal. A “static” manipulation alters a single payoff at a single state (reducing the temptation payoff in the *High* state), and makes simple history-dependent cooperation relatively more attractive. A “dynamic” manipulation alters the transition rule between states to make deviations from joint-cooperation relatively less tempting (holding constant the original pivot’s two stage-games, we make it harder to remain in the *High* state). In both manipulations the direction of the change in behavior is an increase in the selection of efficient outcomes, with less efficient 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, the dynamic externalities are clearly an important factor in the subjects’ behavior. Moreover, while behavior in the treatment with an exogenous but changing state is consistent with the MPE prediction (just under 60 percent have a strong match to it), behavior in the fixed *Low*-state repeated game is much more cooperative and history-dependent. Comparing the two modifications without dynamic externalities, we conclude

that a changing strategic environment does make it harder for subjects to coordinate on history-dependent strategies. But this makes the success for history dependence in the pivot all the more impressive, with its evolving (but endogenous) state.

We also conduct treatments that remove the static strategic effects, retaining the pivot’s dynamic consequences (the endogenously evolving state). 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. Under both parametrizations we find an increase in the frequency of equilibrium Markov play, and a decrease in the frequency of history dependence. The presence of strong static externalities is therefore also identified as an important factor in the selection of history-dependent play.

The final set of manipulations in the paper 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 history dependence.

Our second state-space manipulation adds two endogenous states (each with their own particular stage games) to the pivot game. 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. With more endogenous states, there is greater variation in subjects’ play, and a greater selection of pessimistic play that starts out by defecting. Coordinating on efficient cooperative outcomes seems to become 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) Behavior is particularly sensitive to changes in the transition rule. iv) Increased complexity in the state space does not on its own lead to Markov play

becoming focal. v) 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 state-conditioned behavior.

Clearly, the larger set of dynamic environment is very rich, and our paper only looks at a small family of games within it. Our aim here is to start documenting which broad features of the environment have strong effects on behavior, so that eventually it might be possible to develop 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 topic and outline some implications of our findings for this larger research agenda, and show that some modified rules of thumb from the repeated games literature can be predictive of selection.

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 $t=1,2,\dots$. Underlying the game is a payoff-relevant state $\theta_t \in \Theta$ evolving according to a commonly known transition rule $\psi : \mathcal{A} \times \Theta \rightarrow \Delta\Theta$, 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 $u_i : \mathcal{A} \times \Theta \rightarrow \mathbb{R}$, dependent on both the chosen action profile a_t and the current state of the game θ_t . Preferences over supergames will be represented by the discounted sum (with parameter δ):

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

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 ($\Theta = \{Low, High\}$), and two available actions, $\mathcal{A}_i = \{Cooperate, Defect\}$. 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 random) 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 $\Gamma = \langle \Theta, \theta_1, u_i, \psi \rangle$ 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 we 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, our design uses a partial strategy method that guarantees data collection for at least five periods within each supergame. The method, which is a modification of the block design (cf. Fréchette and Yuksel 2013) implements $\delta = 0.75$ as follows: At the end of every period, a fair 100-sided die is rolled, the result indicated by Z_t . The first period T for which the number $Z_T > 75$ is the final payment period in the supergame.

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

⁴This 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 be lower with experience and do not interact with treatments.

⁵For 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 (Chandrasekhar and Xandri, 2014), the former works 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".

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

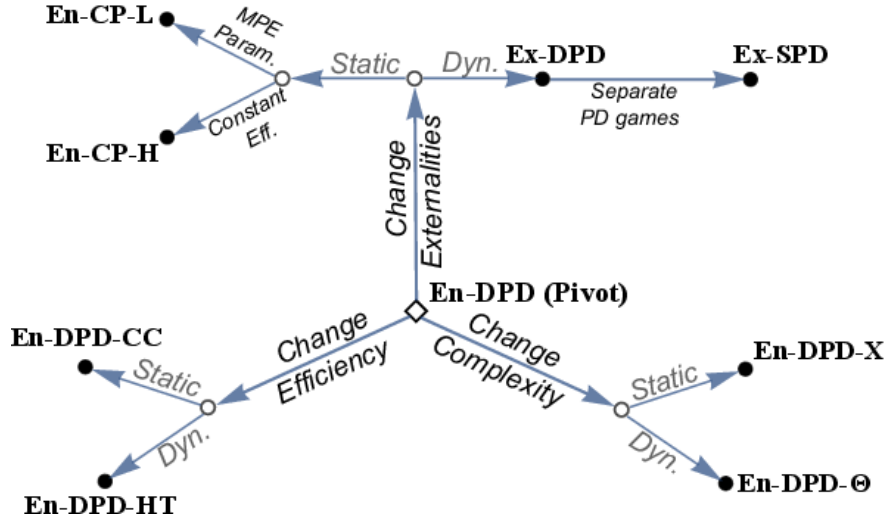


FIGURE 1. Summary of Treatment Design

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 (section 4); ii) changing strategic externalities, how one agent's choice affects the other's payoffs (section 5); and iii) changing the size of the state space (section 6). In each manipulation we change a single feature of our pivot dynamic game, 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 by describing our pivot treatment, and outlining the behavior we find within it, which we do in the next section.

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

TABLE 1. Treatment Summary

Treatment	$ \Theta $	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-CP-L	=	=	=	=	=	=	=	=	=	=
En-CP-H	=	=	=	=	=	=	=	=	=	=

Change State-Space Complexity (Section 6:)

En-DPD- X	22	\div	\div	\div	\div	=	=	\div	\div	\div
En-DPD- $\tilde{\Theta}$	4	\div	\div	\div	\div	=	=	\div	\div	=

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 \div 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.

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

⁸An 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,

TABLE 2. En-DPD

$\theta=\text{Low}$				$\theta=\text{High}$			
		2:				2:	
		C	D			C	D
1:	C	100,100	30, 125	1:	C	200, 200	130, 280
	D	125, 30	60,60		D	280, 130	190, 190

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 strategic externality. The choice of player i in period t affects future states and thus has direct implication on the continuation value for player j .

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

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

A symmetric pure-strategy Markov perfect equilibrium (MPE) is a profile $(M_{\sigma_L \sigma_H}, M_{\sigma_L \sigma_H})$ that is also an SPE of the game. For our pivot there is a unique symmetric MPE, the strategy M_{CD} : both players cooperate in the low state, 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 a SPE. For any sub-game where the game is in the high state, this Markov profile dictates that both agents jointly defect from this point onward, yielding the

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 only transitions back to low following more systemic opportunistic behavior (joint defection).

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 parametrize 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 each state, and is therefore more demanding than Markov strategies such as M_{DD} that choose actions unconditionally; 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 $a_t \neq (C, C)$, 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. 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 yields 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

⁹Expanding 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 the high state. 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.

¹⁰The 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 an SPE).

¹¹We 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 restored; however the payoff difference between efficient and a symmetric solution is small, amounting to 5 cents per player.

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 2) and variation across supergames (Figure 3) while the table (Table 5) provides estimates of the selection frequency for a number of key strategies. While more-detailed regressions are included in the paper’s appendices, to simplify the paper’s exposition we will focus on just these three main sources to discuss our results, with details in footnotes and the appendix.

The first source, Figure 2, presents the most-aggregated information on behavior, the average cooperation rate by state, as well as some basic patterns for behavior within and across supergames. The left-most 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 the low state 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 considerable proportion of subjects are trying to reach more efficient outcomes than the MPE prediction of no cooperation in the high state. Figure 2 also suggests that at least some subjects are not conditioning solely on

¹²Efficient 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).

¹³Every 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 periods 2 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 actions weaken the temptation to deviate.

¹⁴As 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.

the state, that the frequency of cooperation at each state changes as the supergame proceeds. To illustrate this, the figure displays cooperation rates in the first (second) period of each supergame conditional on being in the low (high) state with gray (white) circles. Comparing the placement of each circle relative to the corresponding position of the bar indicates the differences between the initial cooperation rate and the overall average for that state. For En-DPD, the pattern shows much higher initial cooperation levels in the low state, approaching 100 percent in the last five supergames. But the overall low-state cooperation rate is much lower, and if anything falling across the session, suggestive of history-dependent play.¹⁵

While aggregate-level analysis suggests that there is a substantial proportion of subjects who condition their behavior on the history, such results can mask meaningful heterogeneity. For instance, if half of the subjects use strategies that condition on past play, such as S_{CD} , and half use the equilibrium Markov strategies, aggregate cooperation rates may show evidence much more consistent with the latter strategy. To further disaggregate behavior we move to Figure 3. The unit of observation here is the set of choices a subject made within one supergame, which we will refer to as a history. Each history is represented with two numbers: the cooperation rates 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 with bigger circles to represent a greater mass of observations.¹⁶

Figure 3(A) shows that while most histories in the pivot present a perfect or near-perfect cooperation rate in the low state, the dispersion is much larger along the vertical axis, suggesting the presence of three broad categories of cooperation in the high state. The mass of histories near the top-right corner represent supergames where the choices come close to joint cooperation, achievable with 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 choices of subjects 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

¹⁵Table 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.

¹⁶When 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.

¹⁷This method has also been used by Fudenberg et al. (2010), Embrey et al. (2011) and Vespa (2014).

measures the likelihood the observed sequence of choices were produced by each strategy. The method then uses maximum likelihood to estimate a mixture model over the specified strategies (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.

In the reported outcomes in Table 5 we use 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 the efficient outcome: A_{CD} and A_{DD} , where the difference between the two is again on the triggered punishment after a deviation.¹⁹

The SFEM estimates for the pivot treatment in the first column of Table 5 reflect the heterogeneity observed in Figure 3(A). A large mass of behavior is captured by three statistically significant strategies with comparable magnitudes: M_{CD} , S_{CD} and TfT . The frequency of the alternating-state Markov equilibrium 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 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

¹⁸The 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 $\gamma \rightarrow 0$ 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 to fairly well.

¹⁹Efficient 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 $A_{CD} = (A_{CD}^C, A_{CD}^D)$ and $A_{DD} = (A_{DD}^C, A_{DD}^D)$, where A_X^a is the strategy which starts with action a the first time the game enters the high state (see Table 15). 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_{CD}^D and A_{DD}^D .

produces 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.²⁰

3.3. Conclusion. The majority of the data in our pivot is not consistent with the symmetric MPE prediction of joint cooperation in low and joint defection in high. Though we do find close to one fifth of subjects are well matched by the M_{CD} strategy profile, many more attain more-efficient outcomes and 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). Strikingly, the smallest element from the four is the Markov profile M_{CC} indicating that punishments are used often enough for history-dependence to be identified.

Just three strategies account for most of the data in the infinitely repeated PD literature—*Always defect*, the *Grim trigger* and *Tit-for-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. Our findings in a dynamic PD game therefore mirror the results from static PD games. Three strategies account for over 60 percent of the 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 strategy M_{DD} which defects in both states) the vast majority of outcomes and strategies are at or above this level, even where history-dependent punishments are triggered. The MPE strategy is clearly a force within the data, with approximately 40 percent of the estimated strategies using it directly or reverting to it on miscoordination. However, the broader results point to history-dependent play as the norm. The next three sections examine how eight modifications to the pivot alter this finding.

4. CHANGES TO THE EFFICIENT ACTION

Our pivot game is parametrized so the first-best outcomes are asymmetric. Our first set of treatments modify the pivot so that 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: 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.

²⁰We reproduce the SFEM analysis excluding TfT strategies in the appendix (see Table 16). The mass previously captured by TfT is reflected largely through higher estimates for strategies that attempt to jointly cooperate (combined together, S_{CD} , S_{DD} , M_{CC} increase from 42.9 percent to 63.8 percent), while a smaller fraction goes to those that attempt efficient alternation (A_{CD} , A_{DD} together increase from 6.2 percent to 10.8 percent).

4.1. Static Change (En-DPD-CC). Our first modification shifts the efficient action by decreasing the payoff $u_i((D, C), H)$ 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 and 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, the joint payoff in the modification is just 380. Joint cooperation 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.²¹ 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 2 and 3(B), displays many similar patterns (and some important differences) with respect to En-DPD. 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. The high-state cooperation rates for the two treatments are significantly different (at 90 percent confidence), but only for the last five supergames.²² Looking at the supergame level in Figure 3(B), this increase is reflected through much larger concentrations in the top-left corner, perfectly cooperative supergames.

The estimated strategy weights in Table 5 indicate higher frequencies for the 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, as the symmetric trigger with the most weight is the harsher-punishment trigger S_{DD} .

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

4.2. 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 modifies a stage-game payoff in the pivot to make joint cooperation first best, our next treatment

²¹Similar to the reasoning in the pivot, the grim-trigger S_{DD} is not sub-game perfect, as both participants choosing M_{DD} is not an SPE. However, joint cooperation can be implemented with the asymmetric Markov trigger as described in footnote 10.

²²Statistical tests are reported in the appendix Table 9 using a random-effects probit clustering standard errors at the session level.

does this through a change to the transition rule. Exactly retaining the two stage-game payoffs from En-DPD (as given in Table 2), we alter the transition rule in the high-state $\psi(a, H)$ so that any action except for 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 from En-DPD when we alter the dynamics in this way: i) the efficient action in the high state becomes (C, C) as defections yield an inefficient switch back to the low state; 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 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 En-DPD).

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} .²³ Efficient joint cooperation is attainable as an SPE with either symmetric trigger, S_{DD} and S_{CD} .²⁴

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

²³If 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 play 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.

²⁴ TfT is still a Nash equilibrium of the game, but again not an SPE as there is a profitable one-shot deviation along paths that deviate from joint cooperation.

player. Her choice will depend on her belief on the strategy the other will select (M_{CD} or S_{CD}). But the set of beliefs that lead to M_{CD} being a best response in En-DPD is a strict subset of those in En-DPD-HT.

While ex-ante the change in the transition rule might plausibly lead to either more or less Markov play here, the data displays a substantial increase in the selection of efficient outcomes. Looking at the state-conditioned cooperation rates for En-DPD-HT in Figure 2 and comparing it to the pivot, the most apparent results are the significant increase in high-state cooperation.²⁵ Comparing Figures 3(A) and (C) shows a clear upward shift, with the vast majority of supergames 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. Indeed, this focus on harsher punishments for deviations may help drive behavior away from M_{CD} . The set of beliefs that would lead an agent to select M_{CD} is much smaller when she thinks that the other is selecting between M_{CD} and S_{DD} than when the other is selecting between M_{CD} or S_{CD} . We find no evidence for alternating-state Markov play, 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 not consistent with an MPE.²⁶

4.3. Conclusions. In this section we document two modifications to our pivot that examine how changes to the efficient outcome affect subjects' behavior. Though joint cooperation is Pareto dominated by efficient, asymmetric SPE in our pivot, the efficient SPE outcomes in both of our modifications are attainable with symmetric history-dependent SPE. In both treatments, the observed frequency of M_{CD} is negligible, and we observe more instances of successful joint cooperation. The observed move towards efficient outcomes is strongest for En-DPD-HT, when the manipulation affects the pivot game's dynamics.

²⁵The difference is significant at the 99 percent confidence level for the last five supergames.

²⁶The 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 (2014) develops an experimental procedure to obtain extra information that allows to distinguish between such strategies and gain more identifying power.

In both modifications, we find that the rate of MPE use is lower than the pivot, despite a decrease in the attainable efficiency frontier. One way to interpret these treatment effects is that the changes aid subjects' coordination, and reduce strategic uncertainty. In En-DPD, though the payoff differences between efficiency and joint cooperation are small, they may be apparent enough to make coordination more complicated. Our efficiency 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 choice of 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 cooperation. In this section we examine four new treatments, that separate these two types of externality, to see if the strategies subjects select responds to their presence. 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 do the converse: hold constant the pivot's dynamic externalities and remove the static externalities entirely so neither player's choice affects the other's contemporaneous payoff.

5.1. Removing Dynamic Strategic Externalities.

Ex-DPD treatment. To isolate the effects from dynamic externalities in En-DPD our treatments alter the transition rule. The stage-games are identical to those given for the pivot (Table 2) so the static externalities are the same. However, our change to the transition rule removes any interdependence between the current state and the actions chosen last period, and in this way removes dynamic externalities. For our first manipulation of the state-transition rule we choose an exogenous stochastic process for the state's 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.²⁷ Given the exogenous transition rule (Ex-) we label this dynamic PD treatment Ex-DPD.

All MPEs of a dynamic game with an exogenously evolving state are necessarily built-up from Nash profiles in the relevant stage games, as the continuation value of the game is unaffected by current actions (with the strategy's assumed independence). Because the stage-games in each state are PD games this leads to a unique MPE prediction: joint defection in both states. However, other efficient SPE exist that allow for cooperation in the low state and $(C, D)/(D, C)$ alternation in the high state.²⁸

Looking at the experimental results for Ex-DPD, outcomes are starkly different from those where the state's evolution is endogenous. From Figure 2 it is clear that cooperation rates are much lower than the pivot, for both the low and high 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 levels fall from 15 percent at the start of the session, to just under 10 percent in the final five supergames.

The reduced cooperation in Figure 2 is indicated at the supergame-level in Figure 3(D), where the large mass in the bottom-left corner is consistent with sustained defection in both states. This pattern is obviously also reflected in the SFEM estimates reported 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.²⁹

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 over the state helps or hinders cooperation.

Ex-SPD treatment. Our next modification goes further than Ex-DPD, removing the supergame dynamics entirely. To do this we alter the transition rule to keep the game in the same fixed state for

²⁷The Ex-DPD sessions were conducted after the En-DPD sessions were completed. The 60 percent and 80 percent probabilities were chosen to match aggregate outcomes in the En-DPD sessions.

²⁸An 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 from the path. 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.

²⁹We also estimated strategy weights for this treatment adding the history-dependent strategy that supports cooperation only in the low state, described in footnote 28. 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, and S_{DD} is the primary history-dependent strategy.

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 so 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, $\frac{3}{5} \cdot H \oplus \frac{2}{5} \cdot L$. 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), while also assessing subject-level responses to the two different infinitely repeated PD games.

Comparing aggregate-level results in Figure 2 it is clear that cooperation rates in Ex-SPD are higher for both states than for Ex-DPD. Because supergames are in a single fixed state, Figure 3(E) can only show the results on separate axes. The figures shows a large number of supergames with joint defection when the selected state for the supergame is high, but a larger degree of heterogeneity—and more instances of cooperation—when 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 do 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 to the Ex-DPD behavior is starker in the low state. In this case, the frequency attributed to always defect (M_{DD}) is much 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.³⁰

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 games as a repeated PD 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. This finding shows that equilibrium selection does respond to dynamic externalities, suggesting that the endogenously evolving state was a key component in the selection of history-dependent cooperation in the pivot.

³⁰In 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.

TABLE 3. Dynamic Common Pool Treatments

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

$\theta=\text{Low}$

2:

C	D
100,100	100, 125
125,100	125,125

$\theta=\text{High}$

2:

C	D
130,130	130,190
190,130	190 ,190

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

$\theta=\text{Low}$

2:

C	D
100,100	100,125
125,100	125,125

$\theta=\text{High}$

2:

C	D
130,130	130,280
280, 130	280,280

5.2. Removing Static Strategic Externalities. The previous two treatments remove the dynamic externalities, while retaining the pivot's static tensions. We now carry out the converse exercise: turn off the static externalities, while retaining the same dynamic strategic environment. We therefore fix the pivot game's transition rule ψ —joint cooperation is required to transit from low to high, while anything but joint defection keeps the game in high. Our next two treatments instead alter the stage-game payoffs so that each player's static payoff is unaffected by the other player's choice.³¹ Given this restriction, the game can no longer be a PD, and is instead a dynamic common-pool (DCP) problem. In order to make the games comparable to En-DPD, we use two separate parametrizations, with payoffs presented in Table 3.

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 temptation payoff in En-DPD, the payoff from selecting C to that from joint cooperation. Selecting C in the low state therefore allows for the possibility of transiting to high next period if the other also cooperates, but involves giving up on a static gain of 25.

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 from the pivot. The only difference between our two parametrizations is the payoff from choosing D in the high state. In the treatment we will refer to as En-DCP-M the payoff from defecting in high is set to 190, matching the joint-defection payoff in the pivot. In the treatment we will refer to as En-DCP-E the payoff from defection is instead set to 280, matching the temptation payoff in the pivot.

³¹The restriction is therefore that $u_i((a_i, a_{-i}), \theta) = u_i((a_i, a'_{-i}), \theta)$ for all $a_{-i}, a'_{-i} \in \mathcal{A}_{-i}$.

In En-DCP-M the stage-game payoffs match the payoffs attainable with the Markov (hence ‘-M’ prediction from the pivot. The strategy M_{CD} in En-DCP-M yields exactly the same sequence of payoffs (and the same gains/losses from a one-time deviation) as the pivot. Efficient outcomes still involve any combination of $(C, D)/(D, C)$ in the high state, but with this parametrization the payoffs realized from an efficient path are lower than the pivot. In the En-DCP-E treatment the payoffs from any efficient (hence ‘-E’) path are matched to those in the pivot, but conversely payoffs from the most-efficient MPE are higher than in En-DPD.³²

In both DCP treatments the most-efficient pure-strategy MPE uses M_{CD} , though M_{DD} also becomes a symmetric MPE. The efficient outcome in both treatments is the same as in En-DPD, and requires asymmetric play. If coordinated upon, taking turns between cooperation and defection in the high state 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, only A_{DD} is an SPE for En-DPD-E.³³

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

En-DCP-M treatment. The aggregate results in Figure 2 indicate reduced cooperation rates in both states relative to En-DPD. However, the cooperation rate in the low state is now significantly greater than the high state. At the supergame level, Figure 3(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 3(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.

³²If the efficient outcome in the pivot game had involved joint cooperation, we would not have been able to make a clear efficiency comparison to a DCP treatment. Instead, in our larger design, the efficient outcome in the pivot and DCP both require alternation in the high state.

³³Unlike the pivot not all efficient outcomes can be sustained as SPE for the DCP treatments.

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 2), and the largest mass of supergames are in the bottom-right corner of Figure 3(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.

Summary. Our two dynamic common-pool treatments suggest substantial increase in the selection of Markov strategies as we remove static externalities, in particular the MPE. We do find some evidence for greater coordination on the asymmetric efficient SPEs in En-DCP-M relative to the pivot, where the treatment removes second-best symmetric SPE such as S_{CD} . However, as we increase the opportunity costs incurred from initiating the efficient cooperation—giving up 280 instead of 190 by cooperating first—this coordination on asymmetric outcomes disappears. The data indicates that subjects’ strategy selections respond to the presence of static externalities, with a clear increase of symmetric equilibrium Markov play in En-DCP-E relative to En-DPC-M from the pivot.

5.3. Conclusion. Subjects are sensitive to both the dynamic and static strategic externalities in our environments. Whenever we turn either one off we observe an increase in equilibrium Markov play. The absence of endogenous dynamics in the game clearly makes it more difficult for subjects to coordinate on cooperative outcomes. Instead, we observe a substantial increase in the frequency of choices consistent with the MPE, which in the absence of dynamic externalities is M_{DD} (Ex-DPD and Ex-SPD). When we remove the static externalities (the En-DCP treatments) we also find that equilibrium Markov behavior becomes more focal, though in this case captured by the best-case MPE strategy M_{CD} .

In each of the above treatments, removing each type of externality weakens the theoretic desirability of history-dependent strategies. For the treatments without 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 DCP treatments, where there is no static externalities, the change means that efficient outcomes can no longer be symmetric, and constraining to symmetric SPE, M_{CD} is the best outcome. It is possible that the absence of a more-efficient symmetric SPE is the primary driver for the increased Markov play rather than just the removal of static externalities. Though further research will likely separate between these forces more exactly, some evidence already exists. Vespa (2014) 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 mirrors the MPE prediction, suggesting that the absence of static

externalities can be an independent driver for coordination on state-conditioned responses. His game also has a richer state space, and this increased complexity may also contribute to the result. We turn to this potential selection channel in the next section.

6. CHANGES TO THE COMPLEXITY OF THE STATE SPACE

One possible reason for the failure of the MPE to predict subjects' choice in our pivot is that the state space is very simple. We know that 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, 2014). 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 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 set of payoff relevant states in the pivot by adding exogenous, non-persistent shocks, entirely 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 strategic tensions, but with an order-of-magnitude increase in state-space 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.

6.1. Static Complexity (En-DPD-X). One simple way of adding states to the game is through payoff-relevant noise. Our first complexity treatment does this with a commonly known iid payoff shock each period through a uniform draw x_t over the support $X = \{-5, -4, \dots, 4, 5\}$.³⁴ The specific payoffs in each period are given by

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

where $\hat{u}_i(a_t, \theta_t)$ are the En-DPD payoffs given in Table 2. The modification therefore adds an effective shock of $2 \cdot x_t$ (or $4 \cdot x_t$ in the high state) when contemplating a choice between *C* or *D* in low. However, the effect of the shock is static, as the draw next period x_{t+1} is independent with

³⁴This 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).

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 also 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 can be used in a symmetric MPE: cooperate at all states in $\{(L, x) | x \in X\}$, defect for all states in $\{(H, x) | x \in X\}$. The game therefore has the same effective MPE prediction as our pivot.

Moreover, the efficient frontier of the extended state-space game is (for the most part) unaltered, as are the set of simple SPEs.³⁵ Because of the strategic similarity to En-DPD, all the simple symmetric SPE that exist in the pivot have analogs here, while every efficient outcome is again supportable as an SPE by using asymmetric history-dependent strategies. Importantly, given its focality in the pivot, joint cooperation can also be supported here with a Markov trigger.

Examining the results for En-DPD- X in Figure 2, we see qualitatively similar average cooperation levels to those found in the unperturbed pivot. Comparing Figures 3(A) and (H), this similarity extends at the supergame level, though the slightly greater cooperation in both states for En-DPD- X is a little more apparent.³⁶ To make the comparison across treatments cleaner, the SFEM estimates use the same strategies as our previous treatments, and thus ignores strategies that might condition on the shock $x_t \in X$.³⁷ 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 that we note 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 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

³⁵The 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.

³⁶By the last five rounds, the average behavior depicted in Figure 2 for En-DPD- X is significantly more cooperative in both states.

³⁷At the aggregate level, there is evidence of a correlation between the cooperation rate and the value of x in the high state. In the appendix, Figure 4 displays the cooperation rates for different values of x . Table 17 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

		$\theta=\text{Very Low (vL)}$				$\theta=\text{Very High (vH)}$	
		2:				2:	
		C	D			C	D
1:	C	60,60	40,10	1:	C	200,200	10,380
	D	10,40	20,20		D	380,10	85,85

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 our En-DPD- X treatment the added states only affect the current period, and therefore the complexity of conjecture on what might happen that 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 complexity treatment affects the complexity over both what will happen this period (if the new states are reached) and over the conjectures over where the state is headed (at all states). 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- $\tilde{\Theta}$.³⁸

The pivot’s transition rule is modified so that joint cooperation moves the supergame to a higher state (until the ceiling vH is reached), while joint defection moves the supergame to a lower state (until the floor vL is reached). In all other cases the state is held constant. The new transition rule therefore acts as an extension of the pivot’s ensuring 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, as this action profile also transits the game to a better state next period. We calibrate the payoffs so that when both players choose the dominant strategy of C in vL the payoff they receive is matched to the payoff from

³⁸In the experiment the additional states were referred to as the *Green Table* for vL and the *Orange Table* for vH .

joint defection for the low state in En-DPD. For our very-high stage game we choose a PD game (with D as the dominant strategy), choosing the $u_i((C, C), vH)$ payoff so that symmetric joint cooperation in all periods yields an identical sequence of payments in En-DPD and En-DPD- $\tilde{\Theta}$. 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 vH is less efficient than the same alternation in the H state. However, we also 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 between the high and low states superior to alternation between high and very high ($\frac{4}{7} \cdot 190 + \frac{3}{7} \cdot 130 > \frac{4}{7} \cdot 200 + \frac{3}{7} \cdot 85$).³⁹

The pure-strategy MPE of this game is constructed to be 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 this MPE, yielding the same sequence of states and payoffs as M_{CD} in the pivot.⁴⁰ SPE exist that can attain any efficient outcome, while symmetric SPE that maintain joint cooperation do still exist (analogous to S_{CD} and S_{DD}) yielding identical on-path payoffs to those 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 in the pivot for both the low and high states (see Figure 2). The cooperation rates whenever the game reaches the vL state are not shown in the figure, but the average cooperation rates here are high (90 percent for the first five supergames falling to 85 percent for the last five). Additionally, the vH -state's cooperation rates vary from 55 percent in the first five cycles to 60 percent in the last five. Surprisingly, the cooperation rate in the vH state with its more-powerful temptation is actually larger than in the H state.

At the supergame-level, the horizontal axis of Figure 3(I) measures the cooperation rate in either the vL or L states, and the vertical axis cooperation rates in either the H or vH). Given the structure of outcomes in En-DPD- $\tilde{\Theta}$, the top right corner therefore still reflects jointly cooperative outcomes, and the bottom right the equilibrium MPE. Contrasting Figures 3(A) and (I) the observed patterns

³⁹The 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.

⁴⁰Paralleling 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.

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).⁴¹

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.⁴² The equilibrium Markov strategy M_{CD} is still the most popular of the four, but the estimates show an 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} in En-DPD is also consistent with any history-dependent strategy that successfully implements joint cooperation, the increase in M_{DD} in En-DPD- $\tilde{\Theta}$ can be seen as reflecting a small increase in overall Markov play.

History-dependent strategies are similar: 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 which reverts to defection in both the low and high states (and is a good selection when coordination is between conditional cooperation and 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. 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 high state: choosing C is profitable as it keeps the game in high. But since we start both games in the low state, subjects defecting in low never reach the state where the strategy is not sub-game perfect. In the pivot, the first-period cooperation rate is close to 95 percent so that almost all subjects experience the high state, but in En-DPD- $\tilde{\Theta}$ it is significantly lower (at approximately 80 percent).⁴³ Early miscoordination in the low state is particularly damaging: the outcome is statically inefficient, and dynamically it is more likely that the game will cycle between the low and very-low states. In period two, following an initial round where one player defects, only 12 percent

⁴¹For 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- $\tilde{\Theta}$ only 146 of the 210 supergames reach the high state, but 42 percent of these coordinate on joint cooperation.

⁴²For this treatment Table 5 abuses 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 strategies for En-DPD- $\tilde{\Theta}$ is presented in Table 17 of the appendix, where we show that little is lost with the particular restrictions in vL and vH .

⁴³A random-effects probit assessed over the last five cycles rejects equivalence for first-period cooperation in En-DPD and En-DPD- $\tilde{\Theta}$ with 95 percent confidence.

of the 52 En-DPD- $\tilde{\Theta}$ supergames with this history manage to re-coordinate on joint cooperation and reach the high state, while 46 percent have joint defection. In contrast, 55 percent of En-DPD supergames 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, and also less willing to forgive these initial defections. Given the constructed similarities between the two 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.⁴⁴

6.3. Conclusion. The results in the complexity manipulations show that expanding the state-space 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 to succeed. In other words, the presence of a large number of states is not inhibiting subjects from coordinating on more-efficient outcomes than the MPE.

The change in the state-space in En-DPD- $\tilde{\Theta}$ (from 2 to 4 states) is over endogenously reached states, and here we again do not find substantial shifts in the frequency of Markov play relative to the pivot. However, we do find a greater degree of fracture over which strategies subjects are trying to coordinate over, away from joint-cooperation even in the low states, and with a greater frequency of Markov triggers. Though the level of Markov play does not increase *ex ante*, greater miscoordination in supergames leads many to paths of play that are state-dependent and consistent with MPE predictions.

7. DISCUSSION

7.1. Summary of Main Results. Our paper presents experimental results over a core pivot game, and eight modifications 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 tensions within the game, or the dynamics tensions. 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.*

⁴⁴Our 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).

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

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

Evidence: In all treatments with endogenous transitions M_{CD} is the most-efficient MPE prediction. We find that this is the Markov strategy with the highest frequency in En-DPD, En-DCP-M, 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 this strategy is more-likely to be conflated with more-lenient history-dependent strategies.⁴⁵ In treatments with exogenous transitions, M_{DD} is the unique MPE and it is the Markov strategy with the highest frequency.

Result 3 (Coordination and Efficiency). *When the successful implementation of more-efficient outcomes than the best-case MPE requires substantial coordination, we observe higher levels of equilibrium Markov play.*

Evidence: In our environments coordination difficulties present themselves in different ways. First, there may be many outcomes more efficient than the MPE. This is the case, for example, in our pivot En-DPD. The multiplicity is particularly salient because the efficient outcome does not coincide with symmetric joint cooperation. When we alleviate the chances for miscoordination by aligning efficiency with joint cooperation (in En-DPD-CC and En-DPD-HT) we observe much lower levels of MPE play. Second, when the only SPE outcomes with greater efficiency than the MPE are asymmetric (En-DCP-M and EN-DCP-E) we observe even greater levels of MPE play.

Result 4 (Response to Dynamics). *Behavior is very sensitive to how the state evolves.*

Evidence: Removing the endogeneity of the transition (En-DPD→Ex-DPD) yields the most drastic reduction of the cooperation in our treatment set. Moreover, where we remove the dynamics entirely (Ex-DPD→Ex-SPD) cooperation rates increase, suggesting that certainty in the transition is also important. When subjects directly control the evolution of the state, they are generally more successful at coordinating on more-efficient outcomes. Another component of this sensitivity to the dynamics is illustrated in our change to the efficient action. When we change the transition rule (En-DPD→En-DPD-HT) we observe a large shift in both the average behavior, and the frequency of the history-dependent strategies (particularly those with harsher punishments).

⁴⁵As we argued earlier, along the equilibrium path strategies that implement joint cooperation and M_{CC} are identical. Hence, the SFEM cannot separately identify them.

Result 5 (Complexity). *Adding exogenous states (shocks) does not lead to an increase in Markov play, and history-dependent cooperation is still common. However, while adding more endogenous states does not lead to an increase in the selection of Markov strategies, the presence of additional states does alter which strategies are selected, leading to greater miscoordination, and lower cooperation.*

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→En-DPD-X) 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→En-DPD- Θ) overall Markov play stays 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 behavior (see the survey of Dal Bó and Fréchette 2014 for further details). First, the higher the basin of attraction (BoA) of S_D (the grim trigger), the more likely cooperation is to emerge. The basin of attraction for S_D is the set of beliefs on the other player being a conditional cooperator that would make S_D optimal relative to M_D (always defect). In other words, when a relatively low belief on the other cooperating is enough to make conditional cooperation attractive, such strategies are more likely to emerge. Second, cooperative outcomes occur more frequently when cooperation can be supported as an SPE that is risk-dominant. Selecting to cooperate is risk dominant if playing S_D is the best response to the other choosing S_D or M_D with equal probabilities.

While our experiments were designed to investigate behavior across qualitative features of the game, a natural question given our results is whether selection indices like the BoA keep their predictive power in our dynamic environments. Our results point to subjects coordination responding to both the static and dynamic strategic externalities, which motivates extensions to the BoA that integrate both elements (the payoffs and transition rule) into their calculation. In addition, our results also point to symmetric strategies as being focal.

In several of our dynamic games, the direct analogs to the infinitely repeated game strategies (M_{DD} and S_{DD}) are not SPEs, so to evaluate a basin of attraction it is necessary to consider which strategies are reasonable to construct it over. For infinitely repeated PD games, the two strategies compared can be thought of as the MPE (M_D) and a trigger strategy that supports the efficient outcome with a Markov trigger (S_D). For dynamic games we can therefore extend this to define the basin for the best symmetric SPE of the game (for example, in the pivot S_{CD}) relative to the

best symmetric MPE (M_{CD} in the pivot). Hence, we define the basin of attraction for the pivot as the probability that the other chooses the MPE strategy, $p^*(S_{CD}, M_{CD})$, that makes the player indifferent over the two strategies. This index has the interpretation that for any belief that the other player selects M_{CD} lower than $p^*(S_{CD}, M_{CD})$, the best response is to select S_{CD} .⁴⁶

The simple BoA index $p^*(S, M)$ using the best symmetric SPE and MPE is successful in predicting differences in cooperation relative to the pivot for most of our treatments. For our pivot the index is equal to $p^*(S_{CD}, M_{CD}) = 0.754$, so that for all beliefs the other plays M_{CD} from zero to just over three-in-four, it is optimal to choose S_{CD} . This index predicts an increase in history-dependent behavior in En-DPD-CC, which is consistent with our findings, as it increases to $p^*(S_{CD}, M_{CD}) = 0.817$.⁴⁷

In the case of En-DPD- X , the basin $p^*(S_{CD}, M_{CD})$ depends on the realization of the shock x when play is in the high state. If the realization of x is zero, the index coincides with En-DPD. However, the likelihood of coordination changes with x as predicted by the index. For low values of the realization (negative values of x), the index is lower indicating that it cooperation is less likely, while the opposite happens for higher values of x . In fact, as x swings between -5 to 5 the index moves between 0.550 and 0.983 . We find that this is consistent with the findings reported in Figure 4 of the appendix, where the cooperation rate in the high state is positively correlated with x .

The greatest rates of selection for the MPE appear in Ex-DPD, though for this treatment the index has to shift the strategies under consideration. The best symmetric SPE supports cooperation in low and defection in high with an M_{DD} trigger on any deviation (call this strategy X_{DD}), while the MPE prediction here is M_{DD} . The basin of attraction $p^*(X_{DD}, M_{DD}) = 0.255$, and so conditional low-state cooperation is only rational if the belief the other uses M_{DD} is less than one-in-four. Calculating the repeated-game basin $p^*(S_D, M_D)$ for each isolated stage-game in Ex-DPD, we instead have a low-state index of 0.760 and a high-state index of zero (as S_D is not an SPE). Shifts in the cooperation rates from the pivot are matched by changes in the index.⁴⁸

The index is also predictive for the En-DCP treatments. The best-case MPE in these games is the best symmetric SPE, and so this index suggests history-dependent outcomes are not likely, which matches the broad outcomes. However, while the outcome is predicted, the measure ignores some of the important mechanics. The efficient SPE strategies that some subjects are focused on

⁴⁶Defining the BoA more generally. Take any two strategies σ and σ' such that both (σ, σ) and (σ', σ') are SPE the game. The BoA for σ is defined over the belief that the other player uses the strategy $(1-p) \cdot \sigma \oplus p\sigma'$. For some belief p^* , the player is indifferent between playing σ and σ' . While the asymmetric pairs (σ, σ') and (σ', σ) will not be SPE outcomes, it is natural to require that after their initial miscoordination, they subsequently reach a path that is an SPE from that point onward.

⁴⁷The low cooperation in Ex-DPD is also predicted as $\psi_B(\text{Ex-DPD}) = 0.255$, where the strategies considered are a defect-in-High/Grim-trigger in low and M_{DD} , where the Ex-SPD low-state index is 0.760 .

⁴⁸Note that this index also predicts where the coordination occurs. For the pivot it is in period two once the high state is reached. For Ex-DPD and Ex-SPD it is in period one.

this treatments are asymmetric, and within En-DCP-M there is a substantial amount of behavior consistent with such outcomes being selected.

However, we do find a disconnect between the BoA index and the focal strategies in En-DPD-HT. The index predicts *decreased* history dependence in En-DPD-HT relative to the pivot, with a basin of $p^*(S_{CD}, M_{CD}) = 0.634$, but this is the treatment we observe the greatest selection of cooperative, history-dependent strategies. Similarly, in En-DPD- $\tilde{\Theta}$ the index is $p^*(S_{CD}, M_{CD}) = 0.754$, identical to the pivot, but the additional endogenous states lead to greater pessimism at the very start of the game which is therefore not matched by the BoA.

The above simple extension of the basin of attraction to dynamic games offers a rule of thumb selection index, one that integrates both the static and dynamic strategic effects into it. For many of our treatments, more-efficient history-dependent play is predicted over state-dependent play by this index. However, refining this rule presents challenges that future research might address. For instance, it is clear in En-DPD-HT that many subjects use history-dependent conditional-cooperation that punishes with the *worst case* MPE, not the best. Multivariate selection indices that consider coordination across multiple strategies may be more successful. A desirable index might initially compare coordination across S_{CD} , S_{DD} , M_{CD} and M_{DD} in En-DPD-HT, and discern that the $p^*(S_{DD}, M_{DD}) = 0.929$ comparison is the most relevant margin, thereby eliminating the best-case MPE M_{CD} .⁴⁹ Similarly, our complexity treatment with more endogenous states, En-DPD- $\tilde{\Theta}$ shows a decrease in initial cooperation relative to the pivot that is not predicted by the BoA, and further research may help pin down the effects of state complexity on coordination.

8. CONCLUSION

Our paper explores a set of nine dynamic game under an infinite-time horizon in the laboratory. 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 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

⁴⁹One reason for a change in focus might also be changes to the individually rational action in the high state for En-DPD-HT.

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 selected, in the long-run the effect is more pronounced, as many more supergames are miscoordinated. This finding might help explain the greater selection of symmetric MPE in other large state-space experiments on dynamic game (Battaglini et al., 2012, 2014; Vespa, 2014), where increased strategic uncertainty from the many possible future states pushes the game towards simpler Markov 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 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). Or how behavior is affected by the size of the action space: if coordination is important for equilibrium selection, a more complex action-space may play an important role in the selection of strategies.

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.
- Chandrasekhar, Arun G and Juan Pablo Xandri**, “A note on payments in experiments of infinitely repeated games with discounting,” *mimeo*, 2014.
- 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.
- Embrey, Matthew, Guillaume Fréchette, and Ennio Stacchetti**, “Renegotiation in Repeated Games: An Experimental Study,” 2011.
- 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*, 2014.
- 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**, “Cooperation in dynamic games: An experimental investigation,” 2014. UCSB working paper.




FIGURE 2. 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 1 of the supergame for Low states (all supergames), and period 2 for the high state (only those supergames which enter the high state in period 2), except for Ex-SPD, where both circles show period 1 cooperation.

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

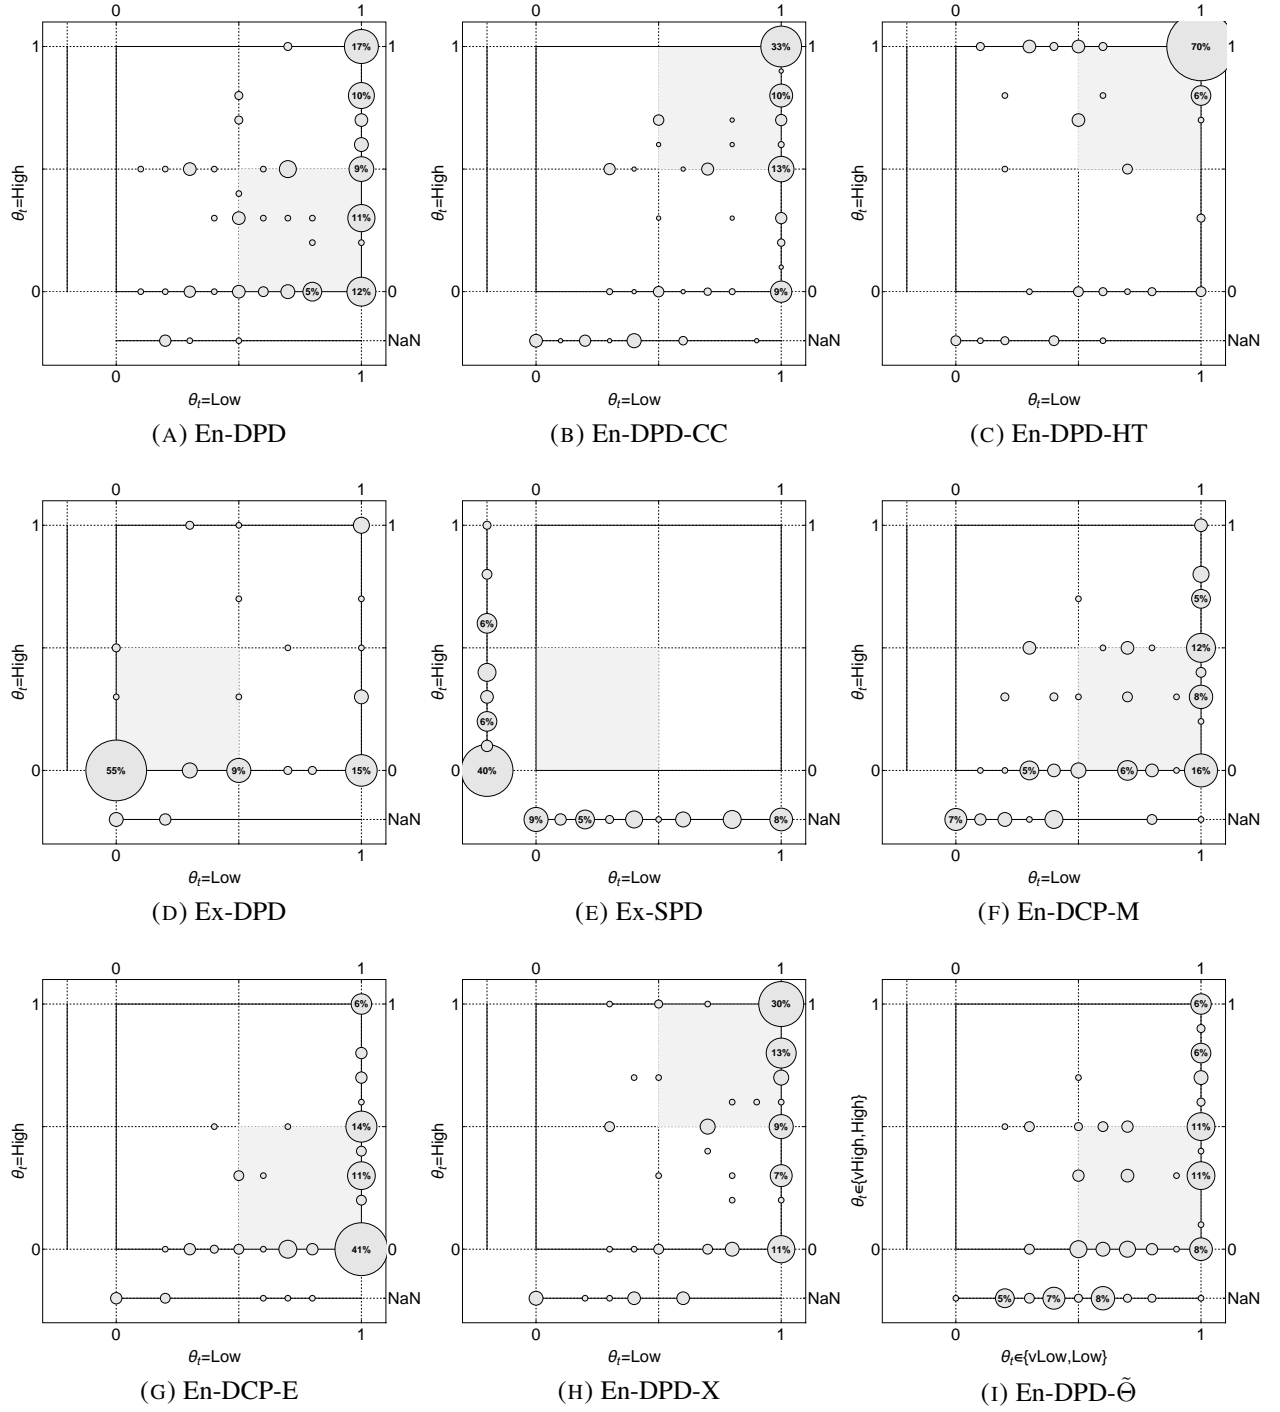


FIGURE 3. Histories (last five supergames)

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

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

		$\theta=\text{Low}$		$\theta=\text{High}$	
		2:		2:	
		C	D	C	D
1:	C	100,100	30, 125	200, 200	130, 250
	D	125, 30	60,60	250, 130	190, 190

TABLE 7. En-DPD-X Stage games

		$\theta=\text{Low}, x$		$\theta=\text{High}, x$	
		2:		2:	
		C	D	C	D
1:	C	$100+x, 100+x$	$30-x, 125+x$	$200+2x, 200+2x$	$130+2x, 280-2x$
	D	$125-x, 30+x$	$60-x, 60-x$	$280-2x, 130+2x$	$190-2x, 190-2x$

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

Treatment	$\theta_t = \text{Low}$			$\theta_t = \text{High}$		
	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)	(***)
En-DPD-HT	0.832	(0.050)		0.979	(0.010)	(***)
Ex-DPD	0.189	(0.059)	(***)	0.012	(0.007)	(***)
Ex-SPD [†]	0.406	(0.062)	(***)	0.079	(0.024)	(***)
En-CP-M	0.638	(0.047)	(***)	0.245	(0.041)	(***)
En-CP-E	0.946	(0.021)	(**)	0.187	(0.047)	(***)
En-DPD-X	0.856	(0.036)	(*)	0.635	(0.055)	(***)
En-DPD- $\tilde{\Theta}$	0.453	(0.033)	(***)	0.254	(0.032)	(***)

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 2 are statistically different from the pivot. The predicted cooperation rates are obtained after estimating a random-effects probit with a dummy variable for cooperation in the left-hand-side, and a constant and a state dummy on the right-hand side.

Table 9 performs a robustness check on the estimates of 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*** (0.028)	0.956*** (0.027)	0.987*** (0.035)	0.971*** (0.026)	0.984*** (0.023)	0.960*** (0.026)	0.963*** (0.025)	0.942*** (0.031)
Treatment	-0.027 (0.048)	0.023 (0.036)	-0.566*** (0.112)	-0.331*** (0.106)	-0.114*** (0.032)	-0.009 (0.040)	-0.042 (0.050)	-0.119* (0.064)
State	-0.331*** (0.057)	-0.333*** (0.057)	-0.335*** (0.057)	-0.336*** (0.056)	-0.329*** (0.059)	-0.332*** (0.057)	-0.330*** (0.058)	-0.331 (0.057)
State \times Treatment	0.136* (0.075)	0.294*** (0.067)	0.067 (0.089)	-0.052 (0.061)	-0.169*** (0.065)	-0.318*** (0.089)	0.163*** (0.059)	0.040 (0.070)
Supergame 12	-0.017 (0.031)	-0.011 (0.023)	-0.015 (0.038)	-0.017 (0.031)	-0.027 (0.026)	-0.031** (0.015)	-0.010 (0.022)	0.003 (0.027)
Supergame 13	-0.024* (0.013)	-0.029* (0.015)	-0.039 (0.033)	-0.069*** (0.026)	-0.074*** (0.027)	-0.011 (0.023)	-0.035 (0.025)	0.019 (0.046)
Supergame 14	-0.017 (0.039)	-0.026 (0.029)	-0.087*** (0.032)	-0.042 (0.028)	-0.076*** (0.020)	-0.042* (0.022)	-0.032 (0.036)	-0.030 (0.027)
Supergame 15	-0.034 (0.021)	-0.024 (0.021)	-0.104* (0.062)	-0.037 (0.032)	-0.055*** (0.017)	-0.026 (0.025)	-0.047*** (0.016)	-0.017 (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 \times Treatment is the interaction of the State and Treatment dummies. Each 'Supergame' variable is a dummy variable that takes value 1 for the corresponding supergame, 0 otherwise. The dependent variable takes value 1 if the subject cooperated, 0 if the subject defected. The data includes all period 1 choices and for all treatments but En-SPD all period 2 choices when the state for that period is High. Each column reports the results of a random effects linear probability model and standard errors (reported between parentheses) are clustered at the session level. Level of Significance: ***-1 percent; **-5 percent; *-10 percent.

TABLE 10. Differences between initial and subsequent period Cooperation Rates

Treatment	$\theta = \text{Low}$			$\theta = \text{High}$		
	$\Delta \Pr \{C\}$	(Std. Err)		$\Delta \Pr \{C\}$	(Std. Err)	
En-DPD	0.498	(0.075)	(***)	0.213	(0.046)	(***)
En-DPD-CC	0.520	(0.066)	(***)	0.135	(0.044)	(***)
En-DPD-HT	0.867	(0.090)	(***)	0.006	(0.014)	
Ex-DPD	0.124	(0.050)	(***)	0.049	(0.022)	(**)
Ex-SPD	0.286	(0.068)	(***)	0.040 [†]	(0.021)	(*)
En-CP-M	0.421	(0.053)	(***)	0.069	(0.039)	(*)
En-CP-E	0.084	(0.039)	(**)	0.040	(0.034)	
En-DPD-X	0.256	(0.065)	(***)	0.256	(0.055)	(***)
En-DPD- $\tilde{\Theta}$	0.421	(0.041)	(***)	0.219	(0.030)	(***)

Note: Figures reflect predicted marginal effect $\Delta \Pr \{C\} = \Pr \{C | \text{Initial Period}, \theta\} - \Pr \{C | \text{Subsequent Period}, \theta\}$ 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 (0.136)	0.175 (0.162)	0.318*** (0.118)	0.045 (0.040)	0.015 (0.062)	0.126* (0.067)	0.391*** (0.129)	0.001 (0.055)
M_{DD}	0.165 (0.110)	0.136 (0.085)	0.000 (0.024)	0.874*** (0.070)	0.407*** (0.144)	0.103 (0.077)	0.041 (0.056)	0.603*** (0.135)
M_{CD}	0.380*** (0.132)	0.135* (0.072)	0.048 (0.030)	0.081 (0.054)	0.432*** (0.110)	0.771*** (0.110)	0.301*** (0.089)	0.396*** (0.124)
M_{DC}	0.262	0.555	0.634	0.000	0.146	0.000	0.267	0.000
γ	1.065*** (0.119)	0.806*** (0.085)	0.408*** (0.070)	0.362*** (0.060)	0.888*** (0.098)	0.615*** (0.133)	0.819*** (0.099)	1.131*** (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-SPD	En-CP-L	En-CP-H	En-DPD- X	En-DPD- $\hat{\Theta}$
<i>LLLL</i>	0.038 (0.004)	0.121 (0.009)	0.057 (0.005)	0.371 (0.023)	0.200 (0.016)	0.048 (0.004)	0.086 (0.008)	0.257 (0.019)
<i>LHLHL</i>	0.038 (0.004)	0.014 (0.001)	0.029 (0.003)	0.00	0.114 (0.010)	0.362 (0.023)	0.019 (0.002)	0.067 (0.006)
<i>LHHHH</i>	0.476 (0.024)	0.607 (0.020)	0.733 (0.019)	0.371 (0.023)	0.19 (0.015)	0.267 (0.019)	0.600 (0.023)	0.314 (0.021)
<i>LHHLL</i>	0.124 (0.011)	0.086 (0.007)	0.010 (0.001)	0.029 (0.003)	0.067 (0.006)	0.01 (0.001)	0.048 (0.004)	0.105 (0.009)
<i>LHHHL</i>	0.105 (0.009)	0.036 (0.003)	0.019 (0.002)	0.038 (0.004)	0.095 (0.008)	0.143 (0.012)	0.067 (0.006)	0.057 (0.005)
<i>LHHHL</i>	0.086 (0.008)	0.064 (0.005)	0.057 (0.005)	0.143 (0.012)	0.076 (0.007)	0.048 (0.004)	0.057 (0.005)	0.095 (0.008)
<i>LHLLL</i>	0.048 (0.004)	0.021 (0.002)	0.067 (0.006)	0.000	0.057 (0.005)	0.010 (0.001)	0.000	0.029 (0.003)
<i>HHHHH</i>				0.629 (0.023)			–	
All Other	0.086 (0.008)	0.05 (0.004)	0.029 (0.003)	0.000	0.200 (0.016)	0.114 (0.010)	0.124 (0.011)	0.076 (0.007)

Note: Data for En-DPD- $\hat{\Theta}$ 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 0.095	CC,CC,CC,CC,CC 0.086	CC,CC,CC,CC,DC 0.057	CC,DC,DC,DC,DC 0.057	
En-DPD-CC	CC,CC,CC,CC,CC 0.264	CC,CC,CC,DC,DD 0.079	CC,CC,CC,CC,DC 0.050	CC,DC,DD,CD,CD 0.043	
En-DPD-HT	CC,CC,CC,CC,CC 0.610	CC,CC,CC,CC,DC 0.124			
Ex-DPD	DD,DD,DD,DD,DD 0.381	DC,DD,DD,DD,DD 0.352			
Ex-SPD (Low)	CC,CC,CC,CC,DC 0.128	DC,DD,DD,DD,DD 0.128	CC,CC,CC,CC,CC 0.128		0.629
Ex-SPD (High)	DD,DD,DD,DD,DD 0.515	DC,DC,DD,DD,DD 0.091	DC,DD,DD,DD,DD 0.076		
En-DCP-M	CC,DC,CD,DC,CD 0.076	CC,DC,DD,CC,DD 0.067	CC,DD,CC,DD,CC 0.057	CC,DD,CC,DD,DC 0.057	DC,DC,DD,DD,DD 0.057
En-DCP-E	CC,DD,CC,DD,CC 0.257	CC,DC,DD,CC,DD 0.105	CC,DC,CD,DC,CD 0.095	CC,DD,CC,DD,DC 0.095	CC,DC,DC,DC,DC 0.048
En-DPD-X	CC,CC,CC,CC,CC 0.162	CC,CC,CC,CC,DC 0.095	CC,DC,DC,DC,DC 0.048		
En-DPD- $\bar{\Theta}$	CC,DD,CC,DD,DC 0.057	DC,DD,CC,DD,CC 0.057	CC,DC,DD,CC,DD 0.048	CC,CC,CC,CC,CC 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_{CD}^C , 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_{CD}^C 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 the SFEM frequencies when we exclude TfT and $STfT$. As argued in footnote 20 this allows us to inquire what proportion actually corresponds to efficient alternation. The estimate for A_{DD} are now positive and significant in En-DPD and En-DCP-M. For En-DPD, this suggests that a mass of approximately 10 percentage points, which is part of the TfT estimate in Table 5 tries to implement the efficient outcome. The estimate is even higher in En-DCP-M.

Table 17 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 4. The figure presents the cooperation rates in low and in high in panels (A) and (B), respectively. Supergames are grouped in blocks of five and the state-space X is divided in three parts: lower than or equal to -3 , between -3 and 3 , and higher than or equal to 3 . Panel (A) shows that the cooperation rate in low is largely unaffected by the choice of x . However, for high state in panel (B) there is a positive effect on cooperation as values of x are higher. Guided by this figure we included two extra strategies in our estimation $M_{CCC,DCC}^x$ and $M_{CCD,DDC}^x$. The supra-script indicates that it is a Markov strategy that conditions on x . The first (last) three values of the subindex indicate the action prescribed in the low (high) state for each of the three elements in the partition of X . Both strategies prescribe the choice of C in the low state for all values of x . This is consistent with the high cooperation rates in panel (A) of Figure 4. In the high state, strategy $M_{CCD,DDC}^x$ prescribes to defect only if the value of x is lower than or equal to -3 , while

$M_{CCC,DDC}^x$ 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 ($S_{CCC,DCC}^x$, $S_{CCC,DDC}^x$).

The estimates in Table 17 are significant only in the case of $M_{CCC,DCC}^x$, 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 17 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* (0.072)	0.177** (0.080)	0.206* (0.122)	0.000 (0.015)	0.104* (0.060)	0.000 (0.014)	0.066 (0.053)	0.092 (0.059)	0.362*** (0.092)	0.016 (0.045)
M_{DD}	0.045 (0.042)	0.087 (0.058)	0.032 (0.041)	0.630*** (0.137)	0.268*** (0.092)	0.733*** (0.118)	0.265** (0.122)	0.069 (0.054)	0.027 (0.041)	0.238** (0.096)
M_{CD}	0.265 (0.102)	0.115 (0.060)	0.050 (0.033)	0.093** (0.047)			0.393 (0.092)	0.682 (0.109)	0.189 (0.071)	0.275*** (0.087)
History-dependent										
S_{DD}				0.194 (0.109)	0.168 (0.100)	0.080 (0.061)				
S_{CD}	0.207** (0.082)	0.194** (0.080)	0.041 (0.078)				0.077 (0.062)	0.088* (0.052)	0.301** (0.122)	0.304*** (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*** (0.061)	0.571*** (0.085)	0.410*** (0.056)	0.357*** (0.042)	0.564*** (0.054)	0.483*** (0.051)	0.805*** (0.094)	0.547*** (0.089)	0.664*** (0.093)	0.775*** (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* (0.071)	0.173** (0.077)	0.068 (0.052)	0.092 (0.058)	0.347*** (0.090)	0.000 (0.040)
M_{DD}	0.024 (0.030)	0.039 (0.039)	0.077 (0.061)	0.048 (0.038)	0.027 (0.034)	0.167* (0.087)
M_{CD}	0.211* (0.127)	0.057 (0.063)	0.311** (0.157)	0.651*** (0.180)	0.138* (0.081)	0.194** (0.090)
M_{DC}	0.000 (0.003)	0.000 (0.010)	0.063 (0.042)	0.000 (0.001)	0.000 (0.014)	0.000 (0.038)
History-dependent						
S_{DD}	0.106 (0.102)	0.227* (0.124)	0.000 (0.042)	0.000 (0.008)	0.069 (0.065)	0.239*** (0.074)
S_{CD}	0.000 (0.076)	0.075 (0.069)	0.000 (0.052)	0.088 (0.056)	0.245** (0.123)	0.045 (0.054)
TfT	0.232*** (0.083)	0.304*** (0.094)	0.134* (0.079)	0.069 (0.078)	0.089 (0.061)	0.068 (0.079)
$STfT$	0.023 (0.027)	0.016 (0.020)	0.000 (0.009)	0.000 (0.002)	0.021 (0.030)	0.000 (0.003)
A_{DD}^D	0.059 (0.039)	0.046 (0.041)	0.161 (0.109)	0.051 (0.050)	0.029 (0.035)	0.099 (0.065)
A_{CD}^D	0.000 (0.103)	0.065 (0.060)	0.080 (0.130)	0.000 (0.143)	0.035 (0.054)	0.000 (0.039)
A_{DD}^C	0.000 (0.058)	0.000 (0.076)	0.044 (0.038)	0.000 (0.002)	0.000 (0.055)	0.063 (0.057)
A_{CD}^C	0.228 (0.058)	0.000 (0.076)	0.061 (0.038)	0.000 (0.002)	0.000 (0.055)	0.126 (0.057)
Parameters						
γ	0.641*** (0.058)	0.529*** (0.071)	0.701*** (0.064)	0.526*** (0.080)	0.635*** (0.088)	0.634*** (0.044)
β	0.826 (0.058)	0.869 (0.071)	0.806 (0.064)	0.870 (0.080)	0.828 (0.088)	0.829 (0.044)

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

TABLE 16. SFEM Output excluding TfT and $STfT$ (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- $\bar{\Theta}$
Markov										
M_{CC}	0.204*** (0.074)	0.268*** (0.095)	0.296** (0.118)	0.000 (0.013)	0.320*** (0.087)	0.519*** (0.112)	0.073 (0.061)	0.100 (0.065)	0.402*** (0.101)	0.029 (0.049)
M_{DD}	0.046 (0.040)	0.056 (0.048)	0.024 (0.017)	0.655*** (0.132)	0.313*** (0.119)	0.168*** (0.046)	0.076 (0.062)	0.048 (0.038)	0.024 (0.034)	0.164* (0.088)
M_{CD}	0.210 (0.132)	0.060 (0.066)	0.040 (0.037)	0.074* (0.043)			0.347** (0.156)	0.647*** (0.176)	0.140* (0.081)	0.170* (0.103)
M_{DC}	0.000 (0.004)	0.000 (0.012)	0.000 (0.027)	0.000 (0.004)			0.063 (0.043)	0.000 (0.002)	0.011 (0.028)	0.000 (0.038)
History-dependent										
S_{DD}	0.153 (0.123)	0.389*** (0.142)	0.616*** (0.110)	0.192* (0.113)	0.367 (0.113)	0.313 (0.113)	0.088 (0.065)	0.000 (0.008)	0.076 (0.077)	0.337** (0.081)
S_{CD}	0.279*** (0.090)	0.117 (0.108)	0.000 (0.055)	0.047 (0.035)			0.104 (0.070)	0.154** (0.074)	0.292*** (0.111)	0.128* (0.068)
A_{DD}	0.108** (0.050)	0.045 (0.048)	0.024 (0.031)	0.033 (0.051)			0.195* (0.110)	0.051 (0.050)	0.028 (0.035)	0.109 (0.081)
A_{CD}	0.000	0.064	0.000	0.000			0.052	0.000	0.026	0.067
γ	0.700*** (0.069)	0.581*** (0.081)	0.368*** (0.043)	0.355*** (0.045)	0.691*** (0.079)	0.771*** (0.204)	0.755*** (0.072)	0.548*** (0.116)	0.656*** (0.090)	0.712*** (0.071)
β	0.807	0.848	0.938	0.943	0.809	0.785	0.790	0.861	0.821	0.803

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


TABLE 17. SFEM Output: Additional Strategies in Complexity Treatments

Strategies	En-DPD-X	En-DPD- $\tilde{\Theta}$
Markov		
$M_{CC} (M_{CCCD})$	0.253*** (0.078)	0.000 (0.041)
$M_{DD} (M_{CDDDD})$	0.027 (0.034)	0.166* (0.088)
$M_{CD} (M_{CCDD})$	0.133* (0.071)	0.176* (0.092)
$M_{DC} (M_{CDDC})$	0.000 (0.013)	0.000 (0.039)
M_{CCCC}		0.059 (0.077)
M_{DDDD}		0.000 (0.038)
$M_{CCC,DCC}^x$	0.203** (0.098)	
$M_{CCC,DDC}^x$	0.002 (0.048)	
History-dependent		
$S_{DD} (S_{CDDDD})$	0.073 (0.062)	0.266*** (0.075)
$S_{CD} (S_{CCDD})$	0.162 (0.119)	0.109* (0.065)
$S_{CCC,DCC}^x$	0.000 (0.019)	
$S_{CCC,DDC}^x$	0.000 (0.020)	
TfT	0.063 (0.056)	0.091* (0.055)
$sTfT$	0.015 (0.024)	0.000 (0.003)
$A_{DD} (A_{CDDDD})$	0.032 (0.036)	0.102 (0.066)
$A_{CD} (A_{CCDD})$	0.038	0.031
γ	0.588*** (0.070)	0.645*** (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 4. 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).