

EXPERIMENTING WITH THE TRANSITION RULE IN DYNAMIC GAMES

EMANUEL VESPA AND ALISTAIR J. WILSON

ABSTRACT. In dynamic environments where the strategic setting evolves across time the specific rule governing the transitions can substantially alter the incentives that agents face. This is particularly true when history-dependent strategies are used. In a laboratory study we examine whether subjects respond to the transition rule and internalize its effects on continuation values. Our main comparison is between an endogenous transition, where future states directly depend on current choices, and exogenous transitions, where the future environment is random and independent of current choices. Our evidence shows that subjects readily internalize the effect of the dynamic game transition rule on their incentives, in line with theoretical predictions.

1. INTRODUCTION

Many economic environments can be modeled with an underlying state variable that changes over time. Combining strategic interaction with an evolving environment, dynamic games provide a broad framework to model economic phenomena. In fact, dynamic games are frequently used in theoretical and empirical applications across virtually every field of economic research.¹ One critical feature in these models is the rule transitioning the state. Holding constant the underlying states and incentives, the transition rule can greatly affect the dynamic incentives players face. In this paper we explore the extent to which human subjects respond and internalize the effects of differing transition rules.

Holding constant the incentives in each state we examine three repeated-game transition rules. Under the endogenous transition rule, the state next period depends on the current state *and* on players' choices. With an endogenous transition, players' relative incentives in the current period are affected not only by the contemporaneous stage-game payoffs, but also by possible future states, where their choices right now can have large effects on available incentives in subsequent periods.

Date: November, 2018.

We would like to thank Gary Charness, Pedro Dal Bó, John Duffy, Matthew Embrey, Ignacio Esponda, Guillaume Fréchette, Drew Fudenberg, John Kagel, Dan Levin, Tom Palfrey, Ryan Oprea and Lise Vesterlund. This research was funded by the National Science Foundation (SES:1629193).

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).

We compare results under an endogenous transition to results under an exogenous transition, where the state next period is independent of players' actions. When future states are independent agents' continuation values become less responsive to history-dependent strategies, as current choices no longer have influence over future payoffs through the state. Finally, to provide a benchmark to translate our results to standard repeated-game environments, we also study a static transition rule where the state is fixed within supergames but varies across them.

Our study isolates the effects of the transition rule on behavior within a dynamic version of an infinitely repeated prisoner's dilemma, which we refer to as a Dynamic Prisoner's Dilemma (DPD). Our DPD extends the standard repeated PD game by adding a single additional state, so that the stage-game can change across rounds. In both states agents face a PD stage game. However, the achievable payoffs in the *high* state are substantially higher than those in the initial *low* state. Our transition rules outlines to participants how the state variable evolves from period two onwards. When transitions are endogenous, both agents need to cooperate in low to transition to high, and once the game enters the high state joint defection is required to move the game back into low.

In standard repeated PD games, history-dependent play allows for the possibility of supporting cooperative play in equilibrium. For instance, trigger strategies that condition cooperation on past play (and punish any deviations with reversion to the stage-game Nash equilibrium) will be subgame perfect equilibria (SPE) of the game so long as players value the future enough. Under our endogenous transition while current choices can still have an implicit effect on future payoffs through history-dependent play, there is an additional direct effect through the stage-game agents play. While joint cooperation is the contemporaneously efficient action in each state it has additional effects on the continuation value, either shifting the state to high or helping to keep it there. Conversely, joint defection not only leads to inefficient present payoffs, it also reduces future values by transitioning the game to the lower-payoff state or keeping it there.

While the stage-game incentives are identical under our exogenous transition rule, the supergame incentives for the continuation game are distinct. Players considering cooperation in the current period still need to consider the present and future payoffs to cooperation across both states, but their choices no longer directly affect the future state. As such the relative incentives for history-dependent cooperation are lower than for the endogenous rule. Finally, in our static treatment the state is selected to be either low or high at the beginning of the supergame, but is fixed within the supergame once chosen. Similar to the exogenous rule, any shift in the continuation value across different actions is now purely driven by the potential for history dependence. However, relative to the exogenous transition, the player no longer has to consider the combination of continuation values at each state within the supergame.

For each of our three transition rules we study two different parameterizations that create variation in the equilibrium sets. In our *easy* parameterization the efficient outcome (joint cooperation in

both states) can be supported as an SPE under all three transition rules. While the endogenous rule does lead to higher continuation values, the best-case equilibrium outcome remains the same in all three treatments. In our *difficult* parameterization, the payoffs in the low state and those resulting from symmetric actions (joint cooperation and joint defection) in the high state are identical to those in *easy*. Holding all other features of the easy games constant, our *difficult* parameterization manipulates the high-state payoffs when agents choose different actions. Holding constant the sum of payoffs to match the *easy* parameterization, the *difficult* treatments increase the temptation to defect in high while simultaneously decreasing the sucker's payoff. The theoretical effect of the change to the difficult parameterization is that sustained joint cooperation across states is now only supportable in equilibrium for our endogenous transition; and not for the static and exogenous rules.

Our experimental results across treatment indicate that subjects strongly react to the transition rule. In the *easy* parameterization, while there are some differences in cooperation rates, the endogenous and exogenous treatments are qualitatively similar. These results show that if the rewards to cooperation are large enough for cooperation to be supported in equilibrium then the precise features of the transition rule have a more-muted effect on final outcomes. However, as we switch to the *difficult* parameterization we detect large treatment effects. Though there are only minimal changes within the endogenous environment, the cooperation rate reductions in the exogenous and static environment are large and significant. These results are in line with the change in equilibrium possibilities for the *difficult* parameterization, where cooperation can still be supported in equilibrium in the endogenous treatment but not in the treatments where the state is drawn independently from choices. This leads to a conclusion that subjects in our experiments internalize the effect of the transition on the dynamic game's continuation value.

Moreover, we also document cross-state effects: difficulties cooperating in the high state lead to reduced cooperation in the low state—even though partial cooperation (cooperating in low, defecting in high) can still be supported with history-dependent play. The clearest evidence for cross-state effects takes place in our static treatments, where the state is assigned at the beginning of the supergame and then fixed. Since our parameterizations only differ with respect to high-state payoffs, static supergames in the low state are theoretically identical in both the easy and difficult treatments. However, low-state cooperation rates are substantially lower in the difficult parameterization. While there is no difference in behavior at the beginning of a session, the difference increases significantly as the session proceeds and subjects gain exposure to the high-state supergame.

Where the results in the static treatments cleanly identify a cross-state learning effect, the evidence from our exogenous treatments identifies a dynamic cross-state effect within the supergame. Many exogenous-transition subject pairs successfully coordinate in the low state early on. However, they find it much harder to sustain this low-state cooperation going forward, significantly less so than

initially successful pairs in the low-state static treatments. One possible channel for this result is strategic uncertainty. While joint-cooperation in the low-state static game resolves much of the uncertainty about a partner's future actions, in the exogenous game uncertainty persists about play in the high state. Despite successful initial coordination in the low state, subsequent defections in the high state contaminate future low-state play. Once returned to the low state, previously cooperative partnerships are now coordinated on defection in both states. In contrast to the results from the endogenous treatments that suggest subjects readily internalize complementarities in the continuation value between the evolving state and history-dependent play, our non-endogenous transition treatments paint a different picture. While future research will help refine this, the results suggest a reduction in long-run cooperation across all states when history-dependent play can only support partial cooperation. The net effect in our treatments is to push long-run behavior towards the history-independent prediction, despite the potential for more-cooperative equilibrium outcomes.

1.1. Literature Review. Our paper is related to a literature on repeated games that evaluates to what extent subjects respond to dynamic incentives. The literature, which was recently surveyed in Dal Bó and Fréchette (2018), has documented that on average subjects do respond to changes in the discount factor and changes in stage-game payoffs as predicted by theory. Specifically, subjects are more likely to cooperate when the future is more valuable (higher discount factor) or when the payoffs to cooperation increase, which is documented also in Dal Bó and Fréchette (2011). Dal Bó (2005) also showed that subjects respond to incentives of the time horizon as predicted by the theory: when the final period is known subjects tend to cooperate less relative to when the ending of the game is determined stochastically. All previous work that we are aware of has kept fixed the role of transition rules. Namely, previous tests of subjects responding to dynamic incentives have kept the transition rule constant to what we refer to as static transition. Our paper shows that subjects also respond to the differential incentives introduced by the transition rules in line with the theory.

There is also a related experimental literature that studies behavior in dynamic games. Most papers focus on issues of equilibrium selection. The set of equilibria in dynamic games can be quite large (Dutta, 1995) and attention in applications is often devoted to symmetric Markov-perfect equilibria (MPE), which are sub-game perfect equilibria (SPE) that do not condition on history. A central question in these papers is to what extent the restriction to MPE is consistent with observed behavior in the laboratory. Clearly, the set of dynamic game environments is very large, however there are some patterns in the literature. For example, Battaglini et al. (2012), Battaglini et al. (2014), Vespa (2017), and Salz and Vespa (2017) study well-known dynamic games with relatively large

state-spaces and find that equilibrium Markov strategies approximate behavior well.² In contrast, the literature on the infinitely repeated prisoner’s dilemma has characterized the conditions under which history-dependent play is likely to prevail. Our two-state DPD can be therefore be thought of as an environment that extends our knowledge on the infinitely repeated prisoner’s dilemma game to two states, where we continue to find evidence consistent with history-dependent play.³ However, beyond equilibrium selection, our study clearly indicates that subjects respond to changes in the transition rule, internalizing its incentive effects on continuations. Understanding that subjects show a clear, theoretically consistent response to the transition rule is a building block for future work in this class of games.

The paper is organized as follows: Section 2 introduces our main treatments, our hypotheses and details of the implementation. We provide our main results in Section 3, where Section 4 summarizes the paper and concludes.

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 $\theta_t \in \Theta$ that starts at some given θ_1 and evolves according to a commonly known (possibly stochastic) 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 the supergame are 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 experiments will examine a family of very simple dynamic environments with an infinite horizon: two players (1 and 2) engage in a symmetric environment with two possible states ($\Theta = \{L(ow), H(igh)\}$) and two available actions, ($\mathcal{A}_i = \{C(operate), D(efect)\}$). Any fewer payoff-relevant states, it is an infinitely repeated game. Any fewer players, it is a dynamic decision problem. Any fewer actions, it is uninteresting.

²For other experimental papers that use dynamic games see Saijo et al. (2014), Benchenkroun et al. (2014), and Kloosterman (2015).

³In a precursor paper to this one, Vespa and Wilson (2017) we expand on the presence of history-dependent and history-independent play in two-state dynamic games. See also Agranov et al. (2018), who also document some conditions under which history-dependent play emerges in a dynamic-game bargaining environment.

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

(A) $\theta=Low$

		2:	
		C	D
1:	C	180, 180	75, 250
	D	250, 75	150, 150

(B) $\theta=High (Easy)$

		2:	
		C	D
1:	C	180, 180	25, 300
	D	300, 25	150, 150

(C) $\theta=High (Difficult)$

FIGURE 1. Stage-Game Payoffs

2.2. **Treatments.** A treatment will be pinned down by the tuple $\Gamma = \langle \theta_1, u_i, \psi \rangle$ indicating 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 states Θ , the set of actions \mathcal{A} , the discount parameter δ and the number of players—will be common.

Endogenous Transitions. We start by describing treatments in which transitioning between states is endogenously determined by the subjects' choices. In period 1, the state is low ($\theta_1 = L$), which means that agents face the stage game in Figure 1(A), where payoffs are in US cents. The next period's state $\theta_{t+1} = \psi(a_t, \theta_t)$ is entirely determined by the actions of the two participants via:

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

This transition rule has a simple intuition: joint cooperation is required to shift the game to the high state from the low state; once there so long the players don't jointly defect the state will remain in high.

Our two treatments with endogenous transitions differ with respect to the game that is played if the state is high. In the *Easy-Endog* treatment the high-state stage game is the one in Figure 1(B), while Figure 1(C) corresponds to the *Difficult-Endog* treatment. Both high-state stage games are parameterized as PD games, though in contrast to the low-state game the returns to symmetric play are much increased. In both high-state stage games the efficient outcome is the same (corresponding to a joint cooperation) and the stage-game Nash equilibrium is the same (corresponding to a joint defection). Where the two high-state parameterizations differ is the payoff from an asymmetric action profile. The *Difficult* game increases the disparity in outcomes between the two participants while holding constant the joint-payoff (\$3.25). As we move from the *Easy-High* game to the *Difficult-High* game the temptation to defect from joint-cooperation is increased by \$0.50 while the sucker's payoff is decreased by \$0.50.

Exogenous Transitions. We have two treatments with exogenous transitions, *Easy-Exog* and *Difficult-Exog*, again varying only over the stage-game payoffs used in the high state. The treatments are identical to the endogenous transition except that the evolution of the state variable is entirely independent of the participants' actions. The state in each non-initial period is determined by the outcome of the lottery $\frac{1}{2} \cdot L \oplus \frac{1}{2} \cdot H$. That is, the game starts in the low state in period one, and from the next period onward the state is high or low with equal chance.

The comparison to the case with endogenous transitions is straightforward. In both cases the state in period one is low, where in period two and forward the state is the result of subjects' choices when the transition is endogenous, and the result of a random process when the transition is exogenous. Clearly, there is a large family of exogenous transition rules that could be used to determine the selection of the next state. We use the simple-to-understand benchmark that makes all states equally likely; where no state has a higher weight a priori.⁴

Conditional on any state, the contemporaneous payoff from any action profile is the same across the endogenous and exogenous transitions. However the two will differ over the game's continuation values. In exogenous, future states are independent of the current action, where endogenous has an explicit dependence on the pair's actions.

Static Transitions

Finally, in our *Static* treatments we shut down the dynamics within the supergame entirely, fixing the state through the transition rule $\theta_{t+1} = \theta_t$. To maintain the same experimental language, we instead allow the starting state θ_1 to vary across supergames. In each supergame we determine the initial state through the lottery $\frac{1}{2} \cdot L \oplus \frac{1}{2} \cdot H$, after which the supergame state is fixed. Again, we have two treatments (*Easy-Static* and *Difficult-Static*) depending on which stage-game payoffs are used for the high state.

The difference between the static and exogenous transitions is again over the continuation value of the game. When transitions are exogenous, even though the state next period is independent the continuation value still involves taking an expectation over future payoffs that takes into account both high and low states. In contrast, with the static transition rule the state next period is known with certainty, so the continuation value does not involve taking expectations over future states.

Summary and Theoretical Properties. Our experiment utilizes a 2×3 -between-subject design over the high-state parameterization (*Easy* or *Difficult*) and the transition rule (*Endogenous*, *Exogenous* or *Static* transitions). Fixing the parameterization—where the variation in the instructions is solely over two payoff numbers—the environment descriptions are identical except for the description and implementation of the transition rule. All else equal, the differing transitions change the future

⁴We will explicitly control for differing state-selection rates in our subsequent analysis, where this choice affords us more power.

incentives within the supergames. As such, the differing theoretical predictions for behavior stem from differences in continuation values. Our experiments' main goal is to evaluate the extent that subjects internalize these differences and respond to variation in the future incentives driven by the transition rule.

In the laboratory we implement all repeated games with a continuation probability of $\delta = 0.75$. Given this common parameter, we now present two properties that hold for both parameterizations and all transition rules. After discussing the common features, we then come back to outline the differences across parameterizations.

A first property of all six dynamic game treatments is that there is a common unique Markov-perfect equilibrium. A symmetric Markov strategy profile is a function $\sigma : \Theta \rightarrow \mathcal{A}_i$, which conditions solely on the current state θ_t , making action choices independent from other elements of the observable supergame history $h_t = \{(a_s, \theta_s)\}_{s=1}^{t-1}$. Given just two states, there are four possible pure-strategy Markov profiles available to each player in our treatments, 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

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

A pure-strategy Markov perfect equilibrium (MPE) is a profile of Markov strategies for both players that is a subgame-perfect equilibrium (SPE) of the game. For each treatment there is a unique MPE with both players selecting M_{DD} , joint defection at every state.⁵ It is straightforward to verify that any other pure-strategy Markov profile cannot be supported in equilibrium in any of the games that we study.⁶

Having the MPE be unique across treatments allows us to rule out differences in observed behavior being attributable to differences (or differential selection) over the MPE set. In addition, it also provides a lower bound on expected payoffs by treatment, as the MPE payoffs in all games coincide with individual rationality.⁷

⁵Abusing notation slightly, wherever we refer to $M_{\sigma_L \sigma_H}$ being an MPE, we will mean that there is a symmetric MPE in which both players use the strategy $M_{\sigma_L \sigma_H}$.

⁶Table 6 in Online Appendix A describes profitable deviations from each other possible MPE. The strategy M_{CD} comes quantitatively closest to being an MPE under the endogenous transition, where we illustrate the calculation. Starting in the low-state the Markov strategy yields the discounted-average payoff of $\pi_L = 4/7 \cdot 100 + 3/7 \cdot 150 = 121.4$ given the predicted alternation across (C, C) in low and (D, D) in high. A one-shot deviation to defect in the low state instead yields $1/4 \cdot 125 + 3/4 \cdot \pi_L$, contradicting the profile being an SPE.

⁷The expected payoff for the MPE is highest under the static transition ($\frac{1}{2} \cdot 60 + \frac{1}{2} \cdot 150$), followed by the exogenous transition ($\frac{5}{8} \cdot 60 + \frac{3}{8} \cdot 150$), followed by endogenous (60).

A second common property is that the joint payoff in every stage-game is maximized with the joint cooperation. This means that in all treatments the Markov strategy M_{CC} implements the efficient outcome, but M_{CC} is not an SPE. When moving from the *easy* to *difficult* parameterization, the design does not change anything for the efficient payoffs. Instead, the change modifies the temptation to defect from the efficient outcome and the cost of miscoordinating on it.

While in all treatments the MPE and efficient path are the same, there are effects from our manipulations on the set of SPE. We focus on a simple history-dependent strategy: a trigger that chooses the efficient Markov profile (M_{CC}) conditional on no observed deviation, but reverts to the MPE profile (M_{DD}) on any defection:

$$S_{DD}^{CC} = \begin{cases} M_{CC} & \text{if no deviation from } M_{CC} \text{ path,} \\ M_{DD} & \text{otherwise.} \end{cases}$$

The S_{DD}^{CC} strategy is the dynamic-game analog to the Grim-Trigger strategy in a repeated PD game. Given that M_{DD} is both an SPE and implements the individually rational payoff, S_{DD}^{CC} provides the best possible chance for efficient play being supported in equilibrium. We now use the two dynamic-game strategies S_{DD}^{CC} and M_{DD} to describe theoretical differences across our experimental design.

In Table 1 we summarize how the incentives are affected by the transition rule and parameterization. Under the *Easy* high-state payoffs given in Figure 1(B), S_{DD}^{CC} is an SPE regardless of the transition rule. In fact, we refer to this parameterization as “Easy” precisely because cooperation can be supported in all three cases, where the transition rule does not qualitatively affect the possibility for supporting efficient play.

While joint cooperation in both states is supportable in equilibrium for our easy parameterization, changing the transition rules might still affect cooperation due to quantitative changes in the incentive to cooperate. To provide a measure for the incentives change, we define $V(X)$ as the discounted-average payoff when both agents use the strategy $X \in \{S_{DD}^{CC}, M_{DD}\}$ from the start of the supergame. We then calculate the relative payoff gains from the cooperative SPE relative to the MPE as $V(S_{DD}^{CC})/V(M_{DD}) - 1$ in the last column of Table 1.⁸ Because this computation is only affected by the main diagonals of the stage-game matrices in Figure 1, the gain only depends on the transition rule, and not the parameterization.

While the gains from cooperation are approximately 35 percent under both the static and exogenous transitions, they are substantially higher (166.7 percent) when the state evolves endogenously. What subjects read in the instructions regarding stage-game payoffs is identical across treatments.

⁸For the static treatment we compute the payoff gain for each possible state first and then weight each payoff by 50 percent, the ex-ante chance of either state being selected at the start of the supergame.

TABLE 1. Treatments

Transition	Param.	S_{DD}^{CC} SPE?	Gain for S_{DD}^{CC} over M_{DD}
<i>Static</i>	<i>Easy</i>	Yes	33.3%
	<i>Diff</i>	No	33.3%
<i>Exogenous</i>	<i>Easy</i>	Yes	38.7%
	<i>Diff.</i>	No	38.7%
<i>Endogenous</i>	<i>Easy</i>	Yes	166.7%
	<i>Diff.</i>	Yes	166.7%

The transition-rule changes affect the way continuation values interact with history-dependent play. If subjects incorporate the quantitative effects of the transition rule on future values, we would therefore expect to observe higher cooperation rates in *Easy-Endog* relative to *Easy-Static* and *Easy-Exog*, even though all three treatments allow for cooperation to be supported in equilibrium.

In contrast to the *Easy* parameterization, when we switch to the *Difficult* high-state parameterization in Figure 1(C), the transition rules effects are not only over the quantitative change in the incentives to cooperate relative to the MPE, but also over a qualitative change in the equilibrium sets. Specifically, S_{DD}^{CC} is an SPE only under the endogenous transition.⁹ When the transition is static or exogenous, cooperation in the high state is no longer supportable in any SPE. Without an interaction between the chosen actions and the state, the shift in the continuation values possible through history dependence is too small to overcome the contemporaneous temptation to defect in the high state. Our *Difficult* parameterizations therefore allow us to examine the extent to which dynamic game behavior that internalizes the transition rule's effects can be predicted by the equilibrium set.

Hypotheses. The above theoretical discussion motivates our main hypotheses. The first hypothesis concerns subjects' ability to internalize the quantitative effect of the transition rule on their future payoff.

Hypothesis 1 (Cooperation Levels). *Ceteris paribus we expect higher levels of cooperation in the Endogenous transition environment.*

This hypothesis is based upon conditional cooperation always being an equilibrium in each of our endogenous transition treatments, and by the quantitatively large payoff gains from cooperation in this treatment. While this first hypothesis considers the cooperation levels across transition rules, our next hypothesis is based on the comparative static response within each transition rule.

⁹It is also an SPE in the *Static* treatments in supergames where the selected state is low. Ex-ante, however, S_{DD}^{CC} is an SPE in any supergame of the *Endogenous* treatments, but not elsewhere.

Hypothesis 2 (Response to Temptation). *Cooperation is more responsive to a change in the high-state temptation under the Exogenous and Static transitions relative to the Endogenous transition.*

Our second hypothesis follows from our variation in the equilibrium set. Cooperation can be supported in equilibrium for the endogenous transition, regardless of the parameterization. However, for the exogenous and static transitions, high-state cooperation is only possible in equilibrium under the *Easy* parameterization. Given some coordination on conditional cooperation for *Easy* across transition rules, our predictions are that the change to *Difficult* is much more deleterious for the static and exogenous transitions than for the endogenous rule.

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

However, subjects are not informed of the outcomes Z_1 to Z_5 until the end of period five. If all of the drawn values are less than or equal to 75, the game continues into period six. If any one of the drawn values is greater than 75, then the subjects' payment for the supergame is the sum of their period payoffs up to the first period T where Z_T exceeds 75. In any period $t \geq 6$, the value Z_t is revealed to subjects directly after the decisions have been made for period t .¹⁰ This method implements the expected payoffs in (1) under risk neutrality. For payment, we randomly select four of the fifteen supergames.¹¹

¹⁰This design is therefore a modification of the block design in Fréchette and Yuksel (2013), in which subjects learn the outcomes Z_t once the block of periods (five in our case) is over. We modify the method and use just one block plus random termination in order to balance two competing forces. On the one hand we would like to observe longer interactions, with a reasonable chance of several transitions between states. On the other, we would like to observe more supergames within a fixed amount of time. Our design helps balance these two forces by guaranteeing at least five choices within each supergame (each supergame is expected to have 5.95 choices). Fréchette and Yuksel (2013) show that "block designs" like ours can lead to changes in behavior around the period when the information on $\{Z_t\}_{t=1}^5$ is revealed. However, such changes in behavior tend to disappear with experience and they show that this does not affect comparative static inferences across treatment.

¹¹Sherstyuk et al. (2013) compare alternative payment schemes in infinitely repeated games in the laboratory. Under a 'cumulative' payment scheme similar to ours subjects are paid for choices in all periods of every repetition, while under the 'last period' payment scheme subjects are paid only for the last period of each supergame. While the latter is applicable under any attitudes towards risk, the former requires risk neutrality. However, Sherstyuk et al. observe no significant difference in behavior conditional on chosen payment scheme, concluding that it "suggests that risk aversion does not play a significant role in simple indefinitely repeated experimental games that are repeated many times."

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.¹² Each session consists of 14 subjects, randomly and anonymously matched together across 15 supergames. We conducted three sessions per treatment, where each session lasted between 70 and 90 minutes with participants receiving average payments of \$19.¹³

3. RESULTS

3.1. Main Comparative Statics. We start our analysis of the experimental results with a comparison of aggregate cooperation rates by treatment, which we use to test our first hypothesis. The first two columns in Table 2 report cooperation rates from the last five supergames in each session.¹⁴ Average cooperation rates are reported for each treatment in our three-by-two design (with subject-clustered standard errors), as well as the difference in the cooperation rate, Δ_Ψ , as we move from the *Easy* to *Difficult* parameterization for a fixed transition Ψ .¹⁵

Inspecting the first three columns in the table, the main takeaway is that cooperation rates in the *Easy* parameterization are similar across the three transition rules, but with large reductions for the exogenous and static transitions in *Difficult*. In contrast, the endogenous transition rule shows a more-muted decrease in response to the increased high-state temptation as we move from the *Easy* to *Difficult* parameterization.

The cooperation rates computed in the first three columns are unweighted, and differences in state composition between transition rules could in principle be driving differing cooperation rates. To control for this we reweight cooperation in our exogenous and static treatments. In the final two columns in Table 2 (again with subject-clustered standard errors) we reweight the cooperation rates to match the exact state mixture observed under the endogenous transitions for each parameterization.

¹²Instructions are provided in Online Appendix B. In the instructions we refer to periods as rounds and to supergames as cycles.

¹³Each subject participated only in one session. All sessions were conducted between January and February of 2018. For each treatment we have three sessions: Session 1, Session 2 and Session 3. To control for the possibility that the particular realization of random numbers may affect our results we proceeded in the following way. The first time we conducted Session X for $X \in \{1, 2, 3\}$ the termination random numbers were selected with a seed set to the time the session started. Subsequent Session X s used the first implementation seed. All Session 1s therefore have the same random termination periods by supergame across treatments, and similarly for Session 2s and Session 3s. We looked for session-specific effects by examining differences across treatments over X . If such differences were significant, there would be evidence that the specific random termination numbers used affected the treatment effects; however, we do not find any evidence of such differences.

¹⁴In addition, we restrict our supergame observations to periods one to five so that longer supergames are not overrepresented.

¹⁵Standard errors are recovered from a subject-clustered linear probability model, where regressors are mutually exclusive treatment dummies, so the method is unbiased.

TABLE 2. Aggregate Cooperation (Last 5 Cycles)

Transition	Unweighted		Δ_Ψ	State-matched	
	Param.			Param.	
	<i>Easy</i>	<i>Diff.</i>		<i>Easy</i>	<i>Diff.</i>
<i>Static</i>	0.604 (0.049)	0.351 (0.044)	-0.252 (0.065)	0.595 (0.050)	0.319 (0.042)
<i>Exogenous</i>	0.561 (0.049)	0.383 (0.047)	-0.177 (0.067)	0.548 (0.050)	0.357 (0.047)
<i>Endogenous</i>	0.633 (0.046)	0.613 (0.035)	-0.020 (0.058)	0.633 (0.046)	0.613 (0.035)

Note: Coefficients in the first two data columns (and standard errors accounting for 252 subject clusters) are recovered from a linear probability model with six treatment-dummy regressors. Coefficients in the state-weighted column are derived from a similar model with the following set of mutually exclusive dummy variables: (i) a treatment dummies for the two *Endogenous* treatments, with coefficients representing $\hat{\Pr}\{C|\Psi_X\}$; and (ii) treatment-state-period dummies for the exogenous treatments, with coefficients representing $\hat{\Pr}\{C|t, \theta, \Psi_X\}$. Reported coefficients for the *Exogenous* and *Static* treatments reflect the weighted sum $Q(\Psi_X)$ across the relevant treatment-state-period coefficient to correct for differing state selection.

Our reweighted averages use the state-period cooperation rates observed under the exogenous and static transitions (where each conditioning variables are mechanically independent in these treatments). Using data from the last five supergames for each *Endogenous* parameterization $X \in \{\text{Easy, Diff.}\}$ we calculate the observed fraction of rounds in the high state in each period t , $\hat{\lambda}_X^t = \hat{\Pr}\{\text{High} | \text{Endog}, X, t\}$.¹⁶ For the transition rule treatment Ψ_X where the state is independently we then estimate the state-period cooperation rate, $\hat{\Pr}\{C|t, \theta, \Psi_X\}$. The state-weighted cooperation rate is then calculated as:

$$Q(\Psi_X) = \frac{1}{5} \sum_{t=1}^5 \left((1 - \hat{\lambda}_X^t) \cdot \hat{\Pr}\{C|t, \text{Low}, \Psi_X\} + \hat{\lambda}_X^t \cdot \hat{\Pr}\{C|t, \text{Low}, \Psi_X\} \right).$$

This reweighted measure therefore maintains as a constant the state composition by period from the endogenous-transition treatment, and allows us to make a controlled comparison across transition rules (fixing the parameterization X).

The main effect of the reweighting is to reduce the effective cooperation rates in our static and exogenous transition treatments. This reduction reflects the greater incidence of the high state under the endogenous transitions, and reduced high-state cooperation under the static and exogenous

¹⁶In the *Easy-Endog* treatment the game is in the high-state 0, 71, 70, 70, and 56 percent of the time, in periods one to five, respectively. For the *Difficult-Endog* treatment the high states is observed 0, 81, 82, 61 and 56 percent, respectively.

rules. However, while there is a change in levels, the qualitative findings remain similar from the unweighted average.

We now make use of the aggregate figures in Table 2 to test our two hypotheses. Hypothesis 1 is that, *ceteris paribus*, the endogenous transition rule produces the greatest cooperation rate. We test this jointly across parameterizations using the vector $\mathbf{Q}(\Psi) = (Q(\Psi_{\text{Easy}}), Q(\Psi_{\text{Diff.}}))$ for each transition rule. Comparing *Endogenous* to both *Static* and *Exogenous* (using Wald tests) we reject the null of no difference with $p < 0.001$. In both parameterizations, we find the greater cooperation under the *Endogenous* transition rule. Looking at the *Easy* parameterization in isolation—rather than jointly with *Difficult*—the endogenous transition is not significantly greater than *Static*, and only marginally significant compared to *Exogenous* ($p = 0.072$ with a one-sided alternative). On the one hand, some of this is a power issue due to our focus on late-session behavior. Expanding our sample to include all cycles instead of just the last five the differences do become significant in the predicted direction ($p = 0.053$ and $p = 0.016$). On the other hand, the small quantitative effects are in line with the equilibrium result that if cooperation can be supported in equilibrium, the magnitude of the payoff gain from cooperating relative to M_{DD} is not of first-order importance.

Result 1: The endogenous transition rule leads to significantly greater cooperation than the static and exogenous transitions. The magnitude of the differences are larger when cooperation can only be supported in equilibrium with the endogenous rule.

Our second hypothesis relates to the reaction to our parameterization shift. Under the assumption that there is some conditional cooperation in our three *Easy* treatments, Hypothesis 2 predicts relative reductions holding constant the transition rule as we shift to the *Difficult* parameterization. Moreover, the relative effects should be largest under the static and exogenous transition rules.

As Hypothesis 2 focuses on the cooperation rate *within* transition, we test this hypothesis by looking at changes to the unweighted cooperation rates across the parameterization, given in the Δ_{Ψ} column in Table 2. Our experimental results indicate a 25 percentage point cooperation reduction (significantly different from zero with $p < 0.001$) for the *Static* transition, and an 18 point reduction in *Exogenous* ($p = 0.009$). In contrast, the reduction in *Endogenous* is quantitatively small (and statistically insignificant) at just 2 percentage points. Testing the reductions across transitions we find the reductions under the exogenous and static rules are significantly different from the reduction under the endogenous rule ($p = 0.024$ using a joint test).¹⁷ We therefore reject the null of no effect in favor of the alternate hypothesis that the endogenous transition rule is less responsive to changes in the temptation to defect.

¹⁷Separately, we find that the reductions in both the *Exogenous* and *Static* treatments are also significantly different ($p = 0.039$ and $p = 0.004$ against one-sided hypotheses).

Result 2: Behavior under the endogenous transition is less responsive to increased temptations to deviate from cooperation than behavior under the exogenous and static transitions, each of which have significant and large reductions.

While we do find directional evidence for Hypothesis 2, there are some potential discrepancies with the theory. Conditional cooperation is *not* possible in equilibrium in both states for the *Difficult-Static* and *Difficult-Exog* treatments. In both environments it is possible to support cooperation in the low state, and so the observed non-zero cooperation could be driven by *Low*-state behavior. However, we find this is not the case. Focusing on behavior solely in the high state, while we do find substantially lower cooperation rates in our *Difficult* parameterization— 21.3 percent high-state cooperation in *Static* and 26.4 percent in *Exogenous*—both of these are far enough from zero that we cannot attribute this to noise.¹⁸ We will come back to this discrepancy in Section 3.3 when we consider learning across supergames and states.

3.2. Cooperation Within a Supergame: History Dependence. The evidence presented so far provides a test of our hypotheses using aggregate measures of cooperation. However, the mechanics of the behavioral predictions can be further scrutinized by studying choices within a supergame. Our theoretical predictions relied upon history dependence, with cooperation being conditioned on past play. A first-order question then is the extent that subjects respond to the history of play. We now use a reduced-form approach to show that the large majority of observed cooperation is indeed history-dependent.¹⁹

The general idea behind our reduced-form evaluation of history dependence is as follows. We focus on the subset of data matching a partial initial history \mathcal{H} motivated by the theory, and study how behavior in the subsequent period depends on the matched player’s behavior. For the *Endogenous* treatments we focus on histories in which both players cooperated in period one (low state) and player i also cooperated in period two (in the high state as both cooperated in low). We therefore focus on supergames with the histories $\tilde{\mathcal{H}} = ((C, C, L), (C, a_2^j, H))$ for either period-two choice of the matched player, $a_2^j \in \{C, D\}$.²⁰ For supergames with this history we then examine the period-three behavior of player i . Specifically, we conduct a regression where the left-hand-side is a dummy reflecting cooperation by subject i in period three, and the right-hand side includes four mutually exclusive indicators: the interaction of a parameterization dummy with a dummy for whether the other player cooperated or defected in period two.

¹⁸These rates are significantly different from the high-state cooperation rates under the *Easy* parameterization, with 56.3 percent cooperation in *Static* ($p < 0.001$) and 45.2 percent cooperation in *Exogenous* ($p = 0.004$).

¹⁹Table 9 in Online Appendix A shows the observed frequencies of the most common histories of play by treatment. The main message from most-common histories is consistent with our reduced-form analysis. When play miscoordinates (one player selects C and the other D) play is typically followed by subsequent punishments.

²⁰Looking at all subject-supergames, this history represents 57.8 percent of our data in the *Easy-Endog* treatment, and 49.1 percent in the *Difficult-Endog* treatment.

TABLE 3. History Dependence

	Static (Coop After $\widehat{\mathcal{H}}$)		Exogenous (Coop After $\widehat{\mathcal{H}}$)		Endogenous (Coop After $\widetilde{\mathcal{H}}$)
	$\theta = \text{Low}$	$\theta = \text{High}$	$\theta_2 = \text{Low}$	$\theta_2 = \text{High}$	$\theta_2 = \text{High}$
Other Defected $_{t-1} \times \text{Easy}$	0.309 (0.073)	0.164 (0.077)	0.222 (0.066)	0.152 (0.050)	0.242 (0.080)
Other Cooperated $_{t-1} \times \text{Easy}$	0.983 (0.013)	0.986 (0.014)	0.945 (0.029)	0.925 (0.032)	0.957 (0.019)
Other Defected $_{t-1} \times \text{Diff.}$	0.247 (0.059)	0.173 (0.063)	0.417 (0.108)	0.314 (0.098)	0.165 (0.044)
Other Cooperated $_{t-1} \times \text{Diff.}$	0.983 (0.012)	1.000 (0.000)	0.992 (0.085)	0.616 (0.083)	0.967 (0.013)
<i>Observations</i>	426	320	417	408	673

Note: Each column represents a separate linear probability model regression where the dependent variable is a dummy indicating cooperation by the subject in period three for Endogenous treatments and in period two in other cases. The right-hand side controls are dummy variables that result from the interaction of a Treatment dummy (where *Easy* and *Difficult* vary depending on the regression as explained below) and a dummy that keep tracks of the other player's behavior in the previous period.

The output of the regression is presented in the last column of Table 3. The likelihood of cooperation decreases in both endogenous transition treatments by more than seventy percentage points if the other player defected in the previous period. In further detail, for *Easy-Endog* the difference is 0.715 ($= 0.957 - 0.242$), while in *Difficult-Endog* it is 0.802 ($= 0.967 - 0.165$).²¹ The reduced cooperation rates in response to the other player's defection is not statistically different across our two parameterizations.²² This result indicates the cooperative response in the *Endogenous* treatments are largely history dependent.

When the transition rule is exogenous and when the state does not change throughout the supergame (*Static*) we instead focus on period two choices of subject i , conditioning on a history where subject i cooperated in period one. That is, we use only the supergames with the histories in $\widehat{\mathcal{H}} = \{(C, a_1^j)\}$, where the first player cooperated in period one, and we will study whether the period-two choice depends on the period-one behavior of the other player (j).²³

²¹The reduction in the likelihood of cooperating in period three is significant for both *Endogenous* treatments (both $p < 0.001$)

²²As these are separate regressions we test the null that the difference .715 is not statistically different from .802, where we find that we cannot reject it ($p = 0.343$). A separate regression on the joint data similarly fails to reject.

²³This history represents 71.9 (70.2) percent of our data in the *Easy-Exog* (*Easy-Static*) treatment, and 59.1 (43.8) percent in the *Difficult-Exog* (*Difficult-Static*) treatment.

In period one of *Exogenous* treatments, all subjects face the low state, but in period two the state is independently assigned. In Table 3 we show the output for two separate regressions, depending on which state was realized for period two. The dependent variable is again a cooperation dummy (though now for period two) with the four controls on the right-hand side again representing whether or not the matched subject cooperated in the previous period interacted with the parameterization. The results are presented in the third and fourth columns for *Exogenous*, and the first two columns for *Static* treatments, in each case conditioning on whether θ_2 was randomly assigned as the *Low* or the *High* state.²⁴

The estimates for the exogenous and static transitions show similarly strong evidence for history-dependent cooperation, at similar levels to the *Endogenous* games. In *Easy-Exog*, if the other player cooperated in period one, the likelihood the subject cooperates in period two increases by more than seventy percentage points regardless of the period-two state.²⁵ Under the *Difficult* parameterization we also find evidence that behavior is conditioned on history, but the effects are smaller than for *Easy-Exog*. In the *Exogenous* treatments if the state is low (high) the likelihood of the subject cooperating increases by 57.5 (30.2) percentage points if the other cooperated in the previous period.²⁶ In low-state *Static* supergames, the difference in the period-two cooperation rate is close to seventy percentage points as we shift the other player's action in the first round. At approximately eighty-percentage points, this effect is slightly higher under the high state, with similar magnitudes across the *Easy*- and *Difficult-Static* treatments.

Though the coordination on cooperative play varies by the transition rule as set out in Results 1 and 2, the above analysis indicates that the large majority of cooperation across treatments is conditional on the supergame history. We summarize the results from Table 3 as:

Result 3: The large majority of cooperative behavior across all of our treatments is history dependent.

3.3. Cross-State Effects.

Hypotheses on Partial Cooperation. Having shown that cooperative behavior is history dependent, we now present more-detailed analyses of behavior across transitions. Recall that differences in the predictions between *Easy* and *Difficult* parameterizations stem from a change in high-state

²⁴While the controls are conceptually similar across regressions in Table 3, the goal of the exercise is not to compare across the three transition rules. In particular, the histories that constrain observations in *Endogenous* are not the same as those in *Exogenous* and *Static*. The goal of the exercise is simply to evaluate whether cooperation within transition is conditional to past play.

²⁵The differences equal .723 ($\theta_2 = Low$) and .773 ($\theta_2 = High$), and each one is statistically significant ($p < 0.001$ in both cases).

²⁶Both figures are relatively high and statistically significant ($p < 0.001$ in both cases). However, if we test whether the difference in the low state is larger than the high state, we can reject the null ($p < 0.001$).

payoffs. We use the name *Difficult* for treatments where our high-state parameterization makes the problem of supporting cooperation more difficult. In principle, making it hard to support cooperation in the high-state also makes cooperation more difficult to achieve in the low state; however, *partial* cooperation remains possible in equilibrium. Specifically, in every treatment the following dynamic-game strategy is a symmetric SPE:

$$S_{DD}^{CD} = \begin{cases} M_{CD} & \text{if no defection in low,} \\ M_{DD} & \text{otherwise.} \end{cases}$$

This strategy can be used to support cooperation in the low state under all three transition rules; where the strategy is unaffected by the parameterization, as the it calls for both parties to defect in the high state. While the quantitative gains for conditional cooperation in the low-state only are smaller, the relative risks from attempting this coordination are constant across parameterizations. Similar to our main hypotheses, we quantify the gains relative to the MPE for partial conditional cooperation as $V(S_{DD}^{CD})/V(M_{DD})-1$, a 25.9 percent and 19.0 percent gain for the exogenous and static transitions, respectively.²⁷

In cases where agents fail to cooperate both in the high state and the low state we will say that there is a cross-state effect. That is, difficulties in supporting cooperation in the high state have translated into difficulties supporting cooperation at all.

Given the reduced gains from *any* cooperation if cooperation in both states is not possible, we formulate the following hypothesis:

Hypothesis 3 (Cross-state effects). *An inability to support cooperation in the high state reduces low-state cooperation.*

The above hypothesis considers the effects on initially cooperative play in the low state, under the idea that the belief that subjects will not cooperate in the high state reduces coordination on low-state cooperation. Additionally, cross-state effects may depend on the transition rule. While a pair might successfully coordinate on cooperative play in the low state under the exogenous transition, miscoordination once the high state is reached could trigger defections, even when the game returns to the low state. In contrast any static game with successful low-state coordination will mechanically be trapped in the low state for the rest of the supergame. As such, successful cooperation in the first period is easier to replicate in subsequent periods.

Hypothesis 4 (Ongoing Cooperation). *Sustaining cooperation in the low state is harder under a exogenous transition where the state changes across the supergame than the static transition with a fixed state.*

TABLE 4. Joint Cooperation in the Low State

Transition	Initial		Ongoing*	
	Easy	Diff.	Easy	Diff.
Static (Low)	0.547 (.064)	*** 0.317 (.060)	0.931 (.075)	0.789 (.093)
Exog. (Low)	0.533 (.046)	** 0.400 (.0486)	0.791 (.082)	*** 0.429 (.077)
Endogenous	0.714 (.046)	0.809 (.046)	—	—

Note: Initial: For each treatment columns show the frequency of pairs of subjects who jointly cooperated in the first low-state round round, standard errors in parentheses drawn from a single OLS regression of a joint-cooperation dummy on the six treatment dummies. *Ongoing*:* Shows the frequency of pairs who jointly cooperated in the fifth round *conditional* on both the low state in round five and the pair having jointly cooperated in round one. Standard errors in parentheses drawn from a single OLS regression of a joint cooperation dummy on treatment dummies for $\{\text{Easy, Diff.}\} \times \{\text{Static, Exog.}\}$, where we do not report figures of Ongoing* for the *Endogenous* treatments, as the low state conditioning is not independent of pair behavior.

Low-State Cooperation Results. Table 4 presents the rates of joint cooperation in the low state using data from the last five supergames.²⁸ In the first two columns we present the joint-cooperation rate in the first low-state round in each supergame, under each separate parameterization. As we move from the easy to the difficult parameterization when the transition rule is endogenous, the joint cooperation rate actually increases, though the effect is very small and not significant.²⁹

In contrast, for the static and exogenous treatments, initial cooperation in the low-state is significantly reduced. Where approximately 55 percent of supergames have joint cooperation under the easy parameterization, we find this falls to 40 and 32 percent in the difficult versions of exogenous and static, respectively. Both reductions are individually significant at the 95 percent level. If we test whether *Easy-Static* together with *Easy-Exog* are different from *Difficult-Static* together with *Difficult-Exog*, the joint test strongly rejects a null of no effect across parameterizations in the two treatments ($p = 0.007$). The data therefore indicates a significant cross-state effect: cooperation not being supportable in equilibrium in the high state significantly reduces low state cooperation.

In the final two columns of Table 4, we present a measure for the rate of ongoing cooperation in the low state. Looking only at those supergames where both parties successfully cooperated in the first

²⁷The endogenous rule provides a 102.3 percent payoff gain from the strategy S_{DD}^{CD} relative to M_{DD} .

²⁸The main findings are qualitatively the same with greater statistical confidence if we use all supergames, and if we examine subject-level cooperation.

²⁹For inference we conduct a common regression with the dependent variable being a supergame-level indication of both participants cooperation on treatment dummies, where we restrict the static treatment to be in the low state.

round, we measure the fraction that are still jointly cooperative in the low state in round five.³⁰ For static supergames that repeat the low-state stage-game, jointly cooperating in the first round leads to a very high likelihood of being jointly cooperative in round five. Approximately 90 (80) percent of the easy (difficult) repeated games are persistently cooperative, and the difference between easy and difficult is not significant. In contrast, while the *Easy-Exog* game has 80 percent of the initially cooperative supergames continuing to jointly cooperate in round five, this falls to a little over 40 percent for the *Difficult-Exog* treatment. This large reduction is statistically significant.

We summarize this evidence as:

Result 4: The evidence is consistent with Hypotheses 3 and 4. The inability to support high-state cooperation has a cross-effect on low-state cooperation, both in initial rates and in the dynamic response.

One strange feature outlined in the evidence for Result 4 is that there are large differences in the static low-state game across parameterizations. Table 4 makes clear that there is a significant reduction in initial low-state cooperation with the static transition rule across parameterizations, with the joint cooperation rate at 55 and 32 percent in *Easy-Static* and *Difficult-Static*, respectively.³¹ But conditional on a low state being initially selected, as the parameter shift only takes place in the high-state payoffs, the supergames are structurally identical. We now show that this is a learned response, and that the reduced cooperation in the high-state games has a contagious effect across supergames.

To start with, in Table 5 we present data on the very first decision subjects make. The table provides the cooperation rates by treatment in the low state, with standard errors and inference from a regression of a treatment dummy for subject cooperation on treatment dummies. We constrain our attention to *Static* treatments first. Notice that while Table 4 shows more low-state cooperation towards the end of each session in *Easy-Static* than for *Difficult-Static*, there is essentially no difference at the beginning of the session. Per Table 5, the initial cooperation rate in the low-state *Static* supergame is actually higher in the difficult parameterization—64 percent compared to 50 percent in Easy, though the difference is not significant. We now show that differential response at the end of sessions is driven by subjects' experiencing the high-state supergame.

³⁰On round 5 we condition on the pair of subjects having successfully cooperated in the first round, and the game being in the low state. We focus on round 5 behavior because this is the last round of the block design for which we have data for all supergames. We do not present results for the *Endogenous* treatments, as being in the low state is not independent of behavior, as it would necessitate joint defections to bring the game back down into the low state. Conditioning on joint cooperation in the low-state, 53 percent of *Easy-Endog* supergames are jointly cooperative in the high state in supergame round five. This compares to 32 percent of *Difficult-Endog* supergames, with the reduction highly significant ($p = 0.006$).

³¹Initial cooperation rates at the subject level are provided in Table 7 of Online Appendix A.

TABLE 5. Initial Cooperation (First Supergame, $t = 1$, $\theta = \text{Low}$)

State	Easy	Diff.
<i>Static</i>	0.500 (0.086)	0.643 (0.126)
<i>Exogenous</i>	0.571 (0.073)	0.381 (0.073)
<i>Endogenous</i>	0.762 (0.073)	0.762 (0.073)

Note: Standard errors in parentheses drawn from a single OLS regression of a cooperation dummy on an exhaustive set of treatment dummies, where we restrict observation of the static treatments to the low state.

Using subject variation in the frequency of exposure to the high-state static supergame in the first ten supergames we can identify the cross-state effect on cooperation. As a simple specification, we examine the decision to initially cooperate in low-state supergames at the end of the session (the last five). We regress these decisions to cooperate on a pair of parameterization dummies and a (standardized) measure of the subject’s exposure to the high-state supergame.³² The estimated regression equation (with subject-clustered standard errors below) is:³³

$$\hat{\Pr}\{a_1^i = \text{Coop} \mid \text{Low}\} = \frac{0.733}{(0.057)} \cdot \delta_{\text{Easy}}^i + \frac{0.563}{(0.076)} \cdot \delta_{\text{Diff.}}^i - \frac{0.101}{(0.048)} \cdot \text{Exposure}_i.$$

The estimated equation indicates a 10 percentage point reduction in cooperation for each standard deviation increase in exposure to the high-state game. The significant reduction in cooperation from exposure to the more difficult supergame ($p = 0.036$) indicates that subjects’ long-run behavior in the repeated game does not treat the supergame environments in isolation. Instead the selection of cooperative outcomes responds to features of the session environment.³⁴

The information in Table 5 also shows that there is a difference between endogenous and the other transition rules, even from the beginning of the session. The *endogenous* treatments have

³²Theoretically, the number of high-state supergames in the first 10 is a Binomial(10, 1/2) random variable. We therefore standardize our exposure variable as:

$$\text{Exposure}_i = \frac{\# \text{ High-state supergames}_i - 5}{\sqrt{2.5}}.$$

³³Unlike our previous models, the linear probability model here is misspecified due to the non-binary exposure variable. However, running a probit and computing marginal effects leads to quantitatively and inferentially equivalent results.

³⁴Looking out-of-sample we find further evidence for Result 6. Using the Dal Bó and Fréchette (2018) meta-study of repeated PD games to make predictions in our *static* transition setting we find less cooperation in our low-state games and greater high-state cooperation. The end effect looks more like a convex combination of the two separate predictions.

significantly greater cooperation from the very first decision, with no effect on initial cooperation from the parameterization.³⁵ It seems that at least some subjects internalize the effects of the endogenous transition rule simply through a description of the environment.

We summarize the findings as:

Result 5: While subjects do internalize the potential gains from an endogenous transition, they do not otherwise display significant differences in initial low-state cooperation. Across the session however, there are significant responses, both to the dynamic environment, and across environments. In *Static* treatments, the more difficult the coordination problem is in high, the lower is the observed cooperation in low.

4. CONCLUSION

Dynamic games are used extensively in theoretical and empirical applications, allowing economists to model and understand strategic environments that evolve over time. The transition rule governing the evolution of the state affects agents' incentives in ways that standard theory will respond to through continuation values shifts. In this paper we present experimental evidence that human-subject behavior mirrors shifts in the continuation-value and equilibrium set driven by changes to the transition rule.

Overall, we find substantial initial and ongoing rates of cooperation in a dynamic game where the transition rule is endogenous, providing a complementarity between contemporaneous cooperation and future states. In contrast, when the transition rule is independent—moving the state either within supergames or across them—we find reduced cooperation. Moreover, the differential cooperation we observe across the transition rules is particularly acute when we shift the game's parameterization. Despite more muted effects (both in theory, and in our data) when the transition rule is endogenous, increases to the temptation to deviate from cooperative play cause much larger effects when the state transitions are independent.

While our results show reduced cooperation in the high state as we manipulate the strategic tensions in this state, a detailed comparison of the behavior in the independent transition rule treatments also finds substantial cross-state effects. Despite partial cooperation being possible in equilibrium in all of our treatments—that is, cooperation in some but not all states—when a treatment does not support cooperation in every state, our results indicate reductions in cooperation in all states. The observed contagions are found both in initial play, and in the ability to sustain cooperation in the long run. While our results are therefore optimistic for subjects' ability to internalize strategic

³⁵A joint test finds that the *Endogenous* treatment has significantly greater cooperation than in either the *Static* or *Exogenous* treatments (respectively, $p = 0.051$ and $p < 0.001$). Pooling the *Easy* and *Difficult* data leads to the same inference. In contrast, a joint comparison of the three easy treatments to the three difficult yields a failure to reject equality ($p = 0.237$)

complementarities in the transition rule, there are also some notes of pessimism. Future research can build upon this, but our finding of substantial cross-state contagion indicates that for games with many states, long-run outcomes may be particularly responsive to whether or not cooperation can be supported in the worst-case state. Rather than more partial outcomes supported by history-dependent strategies, our results suggest outcomes move towards the history-independent Markov predictions.

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.
- Agranov, Marina, Christopher Cotton, and Chloe Tergiman**, “The Use of History Dependence in Repeated Bargaining,” 2018. Working paper.
- 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.
- Benchenkroun, Hassan, Jim Engle-Warnick, and Dina Tasneem**, “An Experimental Study of a Common Property Renewable Resource Game in Continuous Time,” *CIRANO-Scientific Publications 2014s-09*, 2014.
- Bergemann, D. and J. Valimaki**, “Dynamic common agency,” *Journal of Economic Theory*, 2003, 111 (1), 23–48.
- Bó, P. Dal**, “Cooperation under the shadow of the future: experimental evidence from infinitely repeated games,” *The American Economic Review*, 2005, 95 (5), 1591–1604.
- Coles, M.G. and D.T. Mortensen**, “Dynamic Monopsonistic Competition and Labor Market Equilibrium,” *mimeo*, 2011.
- Dal Bó, Pedro and Guillaume R Fréchette**, “The evolution of cooperation in infinitely repeated games: Experimental evidence,” *The American Economic Review*, 2011, 101 (1), 411–429.
- _____, and _____, “On the determinants of cooperation in infinitely repeated games: A survey,” *Journal of Economic Literature*, 2018, 56 (1), 60–114.
- Dutta, P.K. and R. Radner**, “Population growth and technological change in a global warming model,” *Economic Theory*, 2006, 29 (2), 251–270.
- Dutta, Prajit K**, “A folk theorem for stochastic games,” *Journal of Economic Theory*, 1995, 66 (1), 1–32.
- Fréchette, Guillaume R and Sevgi Yuksel**, “Infinitely Repeated Games in the Laboratory: Four Perspectives on Discounting and Random Termination,” February 2013. NYU working paper.
- Hörner, J. and L. Samuelson**, “Incentives for Experimenting Agents,” *mimeo*, 2009.
- Kloosterman, A.**, “An Experimental Study of Public Information in Markov Games,” *mimeo*, 2015.
- Laibson, D.**, “Golden Eggs and Hyperbolic Discounting,” *Quarterly Journal of Economics*, 1997, 112 (2), 443–477.
- 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.
- 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,” 2017. UCSB working paper.

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

Vespa, Emanuel, “An Experimental Investigation of Strategies in Dynamic Games,” 2017. UCSB working paper.

——— **and Alistair Wilson**, “Experimenting with Equilibrium Selection in Dynamic Games,” 2017. Working paper.

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

TABLE 6. Unique MPE in Endogenous Game

Player i/j	M_{DD}	M_{CD}	M_{DC}	M_{CC}
M_{DD}	SPE	Player j deviates to D in low	Player j deviates to D in high	Player j deviates to D in low
M_{CD}	–	Either player deviates to D in low	Player i deviates to D in low	Player j deviates to D in low
M_{DC}	–	–	Either player deviates to D in high	Either player deviates to D in high
M_{CC}	–	–	–	Either player deviates to D in high

TABLE 7. Initial Cooperation Rates

State	<i>Easy</i>			<i>Difficult</i>		
	<i>Static</i>	<i>Exog</i>	<i>Endog</i>	<i>Static</i>	<i>Exog</i>	<i>Endog</i>
Initial Cooperation (All Supergames)						
<i>Low</i>	.726 (.049)	.719 (.057)	.822 (.053)	.602 (.061)	.591 (.061)	.841 (.042)
<i>High</i>	.676 (.060)	.552 (.053)	.834 (.054)	.356 (.055)	.316 (.053)	.678 (.062)
Initial Cooperation (Last Five Supergames)						
<i>Low</i>	.736 (.062)	.733 (.063)	.824 (.057)	.558 (.076)	.610 (.069)	.886 (.047)
<i>High</i>	.692 (.067)	.510 (.070)	.800 (0.066)	.300 (.064)	.302 (.063)	.653 (.065)

Note: The initial cooperation rate captures the frequency of C choices in each state using the first choice a subjects make in that state within the supergame. In the case of the low state, only period-one choices are included.

TABLE 8. P-values of hypothesis tests between Initial Cooperation Rates (Last Five Supergames)

	S_L^{Esy}	S_L^{Dif}	S_H^{Esy}	S_H^{Dif}	Ex_L^{Esy}	Ex_L^{Dif}	Ex_H^{Esy}	Ex_H^{Dif}	En_L^{Esy}	En_L^{Dif}	En_H^{Esy}	En_H^{Dif}
S_L^{Esy}	-	.072	.528	.000	.977	.176	.017	.000	.301	.056	.480	.358
S_L^{Dif}	-	-	.187	.001	.078	.619	.642	.010	.006	.000	.017	.345
S_H^{Esy}	-	-	-	.000	.656	.391	.061	.000	.137	.019	.253	.673
S_H^{Dif}	-	-	-	-	.000	.001	.028	.983	.000	.000	.000	.000
Ex_L^{Esy}	-	-	-	-	-	.187	.000	.000	.289	.053	.465	.375
Ex_L^{Dif}	-	-	-	-	-	-	.313	.000	.018	.001	.047	.648
Ex_H^{Esy}	-	-	-	-	-	-	-	.028	.001	.000	.003	.136
Ex_H^{Dif}	-	-	-	-	-	-	-	-	.000	.000	.000	.000
En_L^{Esy}	-	-	-	-	-	-	-	-	-	.405	.784	.050
En_L^{Dif}	-	-	-	-	-	-	-	-	-	-	.290	.004
En_H^{Esy}	-	-	-	-	-	-	-	-	-	-	-	.113
En_H^{Dif}	-	-	-	-	-	-	-	-	-	-	-	-

Notes: To compute these p-values we first run a regression in which the unit of observation is the choice a subject makes in a round of a supergame. The sample is constrained to the last five supergames and to rounds in which the subject makes the first choice in each state. The dependent variable takes value 1 if the subject decided to cooperate and 0 otherwise. The right-hand side includes a fully saturated set of dummies that account for differences in cooperation rates across three dimensions: the treatment (*Easy-Endog*, *Easy-Exog*, *Easy-Static*, *Easy-Endog*, *Easy-Exog*, *Easy-Static*), the state (Low, High). Standard errors are clustered by subject. The table reports the p-values of bilateral comparisons between coefficients for the treatment cross state dummies. The table reports the p-value of a t-test in which the null hypothesis is Row Estimate=Column Estimate. There is one row per (initial-cooperation rate) coefficient estimate and one column per (initial-cooperation rate) coefficient, where notation is as follows. *S*, *Ex*, and *En* capture whether the coefficient corresponds to a Static, Exogenous or Endogenous treatment, respectively. The superscript (*Esy*, *Dif*) identifies if the coefficient corresponds to a Easy or Difficult parameterization, respectively. The subscript (L, H) identifies if the coefficient corresponds to behavior in the Low or High state, respectively. For example, En_L^{Dif} corresponds to the coefficient estimated for *Difficult-Endog* × Low State.

TABLE 9. Common Sequences of Actions as Percent of Histories (Last 5 supergames)

Treatment	5 or more observed Supergames				
<i>Easy-Endog</i>	CC,CC,CC,CC,CC 37.1	DC,DD,DD,DD,DD 10.5	CC, DC,DC,DD ,DC 6.7	DC,DC,DD,DD,DD 4.8	
<i>Diff-Endog</i>	CC,CC,CC,CC,CC 20.0	CC, DC,DD ,DC,DD 6.7	CC,CC,CC,CC,CC 5.7	CC, DC,DD ,DD,DD 4.8	
<i>Easy-Exog</i>	CC,CC,CC,CC,CC 36.2	DC,DD,DD,DD,DD 20.0	DD,DD,DD,DD,DD 5.7		
<i>Diff-Exog</i>	DC,DD,DD,DD,DD 22.9	DD,DD,DD,DD,DD 18.1	CC,CC,CC,CC,CC 18.1	DC,DC,DD,DD,DD 4.8	
<i>Easy-Static (Low)</i>	CC,CC,CC,CC,CC 50.9	DC,DD,DD,DD,DD 18.9	DC,DC,DD,DD,DD 9.4		
<i>Diff-Static (Low)</i>	CC,CC,CC,CC,CC 25.0	DC,DD,DD,DD,DD 20.0	DD,DD,DD,DD,DD 13.3	DC,DC,DD,DD,DD 11.7	
<i>Easy-Static (High)</i>	CC,CC,CC,CC,CC 42.3	DC,DD,DD,DD,DD 25.0	DD,DD,DD,DD,DD 9.6		
<i>Diff-Static (High)</i>	DD,DD,DD,DD,DD 46.7	DC,DD,DD,DD,DD 22.2			

Note: In *Endog* and *Static* treatments high-state action pairs are displayed in bold face.