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 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 actions. Our evidence shows that subjects readily internalize the effect of the dynamic game transition rule on their incentives, in line with history-dependent 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. As such, 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 governing the state transition. Holding constant the underlying states and incentives, the transition rule can greatly affect the dynamic incentives at play. In this paper we explore the extent to which human subjects respond to 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,

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

where their choices right now can have large effects on available incentives in subsequent periods. 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 an 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.

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 periods. In both states agents face a PD stage game. However, the achievable payoffs in the *high* state are substantially better than those in the initial *low* state. Our transition rules outline 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 cooperation in equilibrium. For instance, trigger strategies that condition cooperation on past play (and punish any deviations with reversion to the stage-game Nash) will be sub-game perfect equilibria (SPE) of the game so long as players value the future enough. Under our endogenous transition current choices can still have an implicit effect on future payoffs through history dependence in actions, but there is an additional direct effect through the next period stage game. While joint cooperation is always the contemporaneously efficient action, it provides additional increases to the continuation value under the endogenous transition: shifting the next-period 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.

The stage-game incentives under our exogenous transition rule are identical to those under the endogenous rule, however, the continuation values are distinct. Players considering history-dependent cooperation in the current period still need to consider the present and future payoffs across both states, but the transition rule removes the ability to directly affect future states. As such, the relative incentives for 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 then fixed within the supergame once chosen. Similar to the exogenous rule, any change in the continuation value is purely driven by history dependence. However, relative to the exogenous transition, the player no longer has to consider a combination of state-conditioned continuation values.

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) are 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 parameterization shift is that sustained joint cooperation across states is now only supportable in equilibrium under the endogenous transition; and not for the static and exogenous rules.

Our experimental results across treatments 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 transitions are qualitatively similar. Where the rewards to cooperation are large enough for cooperation to be supported in equilibrium, the precise features of the transition rule seem to 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 observed changes within the endogenous environment, 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 high-state cooperation *is* an SPE action in the endogenous treatment, but not in the treatments where the state is drawn independently. 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 identical in both the *easy* and *difficult* treatments. However, cooperation rates in the identical low-state supergame are substantially reduced as we move to the difficult parameterization. While there is no difference in behavior at the beginning of the session, significant differences emerge 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 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.

Our results from the endogenous treatments suggest subjects readily internalize complementarities between history-dependent play and the transition rule's effects on the continuation value, but our non-endogenous transition treatments paint a different picture. Where future research will help in refining this, our results suggest substantial reductions 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 if there is some state where cooperation cannot be supported, despite the potential for more-cooperative play at other states.

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 (see 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's main contribution is to show that subjects also respond to the differential incentives introduced by the transition rules.

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 the subgame perfect equilibria (SPE) that do not condition on history. A central question in these papers is to what extent the MPE restriction is consistent with observed behavior. 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 $t = 1, 2, ...$ 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 $\psi : \mathcal{A} \times \Theta \to \Delta\Theta$, so that the state next period is given by $\theta_{t+1} = \psi(a_t, \theta_t)$. The preferences for each player i are represented by a period payoff $u_i : \mathcal{A} \times \Theta \to \mathbb{R}$, dependent on both the chosen action profile a_t and the current state of the game θ_t. Preferences over the supergame are represented by the discounted sum (with parameter δ):

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

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, ($\mathcal{A}_i = \{C(ooperate), D(efect)\}$). 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.

---

[2]For other experimental papers that use dynamic games see Saijo et al. (2014), Benchekroun et al. (2014), and Kloosterman (2018).

[3]In 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.

FIGURE 1. Stage-Game Payoffs

2.2. **Treatments.** A treatment will be pinned down by the tuple $\Gamma =< \theta_1, u_i, \psi >$ 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 *Diff-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). The high-state parameterizations differ in the payoffs when one player cooperates and the other defects. Given asymmetric actions the *difficult* game increases the disparity in outcomes 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 *Diff-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 starting from period two 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 treatments will differ in the 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-transition 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 *Diff-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 values. When transitions are exogenous, even though the state next period is independent, calculating the continuation value still involves taking an expectation over the future that takes into account both high- and low-state realizations. 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 3×2-between-subject design over the the transition rule (endogenous, exogenous or static transitions) and the high-state parameterization (easy or difficult). 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

---

[4]We will explicitly control for differing state-selection rates in our subsequent analysis, where the choice of having states be equally likely affords us more power.

the future incentives within the supergames. As such, the differing theoretical predictions for behavior stem from differences in continuation values. Our experiment's 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]

A second common property is that the the joint payoff in every stage-game is maximized with joint cooperation. This means that in all treatments the Markov strategy M_{CC} implements the efficient

---

[5]Abusing 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 $M_{\sigma_L \sigma_H}$.

[6]Table A1 in the Online Supplement 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.

[7]The 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).

outcome, but M_{CC} is not an SPE. When moving from the *easy* to the *difficult* parameterization, the design does not change 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:*C*

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

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 supergame 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. 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

---

[8]For 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?	S_{DD}^{CC} gain over M_{DD}
Static	*Easy*	**Yes**	33.3%
	Diff	No	33.3%
Exogenous	*Easy*	**Yes**	38.7%
	Diff.	No	38.7%
Endogenous	*Easy*	**Yes**	166.7%
	Diff.	**Yes**	166.7%

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-rule effects are not only over the quantitative change in the incentives to cooperate relative to the MPE, but also a qualitative change in the equilibrium sets. Specifically, S_{DD}^{CC} is an SPE only under the endogenous transition. 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.

Hypothesis 2 (Response to Temptation)**.** *Cooperation is more responsive to a change in the high-state temptation under the exogenous and static transitions rule relative to the endogenous rule.*

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 transition rules than for the endogenous rule.

The analysis so far focuses on two extreme strategies (S_{DD}^{CC}, which would support cooperation in every state and M_{DD}, which captures individual rationality), but partial cooperation is also feasible. Notice first that because S_{DD}^{CC} is an SPE in *easy* and in the *Diff-Endog* treatments, the maximal theoretical rate of cooperation that can be supported in an SPE for these treatments is 100 percent. While full cooperation is not supportable in *Diff-Exog* and *Diff-Static*, cooperation in one of the two states is possible. Specifically, there is an SPE which supports cooperation in the low state only, which means that the maximal theoretical cooperation rate for these treatments is 100 percent in the low state and 0 percent in the high state.

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.[9] This method

---

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

implements the expected payoffs in (1) under risk neutrality. For payment, we randomly select four of the fifteen supergames.[10]

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.[11] Each session consists of 14 subjects, randomly and anonymously matched together across 15 supergames. We conducted three sessions per treatment (leading to a total of 252 subjects) where each session lasted between 70 and 90 minutes with participants receiving average payments of $19.[12]

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 two hypotheses. Each column pair in Table 2 reports unilateral cooperation rates from the last-five supergames of our experimental sessions—one column for each parameterization, where we break out the three transition rules by row.[13] The column pairs present unilateral cooperation rates with subject-clustered standard errors, first aggregated across states (overall cooperation), and then conditional on the state (state-conditioned cooperation).[14]

We provide two measures of unconditional cooperation. The first measure computes the rate of unilaterally cooperative choices in the raw unweighted data. Our second measure is constructed to make a fair comparison across transition rules, as differences in state composition across treatments

---

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

[11]Instructions are provided in the supplementary material. In the instructions we refer to periods as rounds and to supergames as cycles.

[12]Each 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 Xs 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.

[13]In addition, we restrict supergame observations to periods one to five so that longer supergames are not over-represented.

[14]Standard 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 Unilateral Cooperation (Last 5 Supergames)

| Transition | Overall | | | | | State-Conditioned | | | |
| | Unweighted | | | Weighted | | Low State | | High State | |
	Easy	*Diff.*	Δ	*Easy*	*Diff.*	*Easy*	*Diff.*	*Easy*	*Diff.*
Static	0.604	0.351	-0.252	0.595	0.319	0.643	0.455	0.563	0.213
	(0.049)	(0.044)	(0.065)	(0.050)	(0.042)	(0.056)	(0.058)	(0.060)	(0.045)
Exogenous	0.561	0.383	-0.177	0.548	0.357	0.610	0.470	0.483	0.263
	(0.049)	(0.047)	(0.067)	(0.050)	(0.047)	(0.050)	(0.052)	(0.056)	(0.051)
Endogenous	0.633	0.613	-0.020	0.633	0.613	0.484	0.620	0.763	0.608
	(0.046)	(0.035)	(0.058)	(0.046)	(0.035)	(0.063)	(0.048)	(0.042)	(0.043)

Note: Coefficients and subject clustered standard errors in the data column pairs (Unweighted/Low/High) are recovered from a linear probability model with six treatment-dummy regressors on the relevant subsamples. Weighted coefficients for the exogenous- and static-transition treatments reflect the weighted sum of the treatment-state-period coefficients to match the state-period composition in the relevant endogenous treatment, where standard errors are recovered through a linear combination of state-period-treatment dummies. The Δ column reflects the difference in raw cooperation rates for the relevant transition rule.

could be driving differences in raw cooperation rates. To control for this we construct a *weighted* measure of cooperation that exactly matches the state composition to the endogenous transition (by parameterization). We then use the realized weights to compute a state-weighted cooperation measure for our exogenous and static treatments that has the exact same state-composition as the relevant endogenous treatment.[15]

The data presented in Table 2 provide strong support for Hypothesis 1, which we summarize as our first result:

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

Hypothesis 1 predicts that, ceteris paribus, the endogenous transition rule produces the greatest cooperation rate. Aggregating across parameterizations, the cooperation rate in treatments with the endogenous rule is 62 percent, which we contrast to 48 and 47 percent under the static and

---

[15]For the weighted columns we use data from the last-five supergames in each *Endogenous* parameterization $X \in$ {Easy, Diff.} to calculate the fraction of high-state periods in period t, $\hat{\lambda}_X^t = \hat{\Pr}\{\text{High} \,|\text{Endog}, X, t|\}$. We then assemble the Weighted unilateral cooperation rates for Static and Exogenous as:

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

Because each parameterization X has a different weighting, we therefore do not make relative-effect comparisons *within* transition on this column in the table to avoid confusion.

exogenous transitions. Comparing the endogenous transition to both the static and exogenous transitions (using Wald tests jointly over both parameterizations) we reject the null of no difference with $p < 0.001$. While we find greater cooperation in both endogenous rule treatments, when we examine the *easy* parameterization in isolation, the endogenous transition is not significantly greater than the static or exogenous rules. On the one hand, some of this is due to a loss of power given our focus on late-session behavior. Expanding our sample to include all supergames (instead of the last five) the differences do become significant in the predicted direction ($p = 0.084$ and $p = 0.030$ against one-sided alternatives). 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.

While the aggregate evidence supports the comparative-static predictions across treatments, the *state-conditioned* cooperation rates reported in Table 2 help in evaluating how far behavior is from the theoretically maximal cooperation rates.[16] In the *Easy-* and *Diff-Endog* treatments theory allows for a 100 percent cooperation rate in both states. However, in both cases behavior is well below this prediction. The *Easy-Endog* treatment in the high state comes closest to the prediction, with 76 percent cooperation. Similarly, in all of the exogenous and static transition-rule treatments it is theoretically possible to support 100 percent cooperation in the low state. However, we observe much lower cooperation rates, closer to fifty percent on average. As we will document below (section 3.2) one reason that cooperation rates are not 100 percent is that behavior is both history dependent and imperfectly coordinated, where punishments are enacted given a deviation from cooperative play. Finally, we also point out that while it is not possible to support cooperation in the high state for our *Diff-Exog* and *Diff-Static* treatments, we observe cooperation rates close to 25 percent. Consistent with the relative prediction, these are the lowest cooperation rates in the table, but it does indicate a non-negligible number of subjects do try to cooperate even when it is not theoretically supportable.

Evidence from the state-conditioned cooperation rates again provides clear support for Hypothesis 1. However, we note that comparing state-conditioned rates across transitions can be misleading. For example, consider conditioning only on the high-state periods. This will mechanically select pairs that managed to support cooperation under the endogenous transitions, but will not create any selection for the exogenous and static transitions. To make a fair comparison across transitions—and control for the induced state composition in the exogenous and static treatments—we construct a weighted cooperation rate that matches the state composition to the relevant endogenous treatment. This leads to a cooperation measure that is directly comparable across transition

---

[16]The high-state incidence is 55.2 and 52.1 percent in the *Easy-Endog and Diff-Endog* treatments across the first-five supergame periods, respectively. In contrast it is, respectively, 40.1 and 39.9 percent in the *Easy-* and *Diff-Exog* treatments (40 percent induced), and 49.5 and 48.6 percent in the *Easy-* and *Diff-Static* treatments (50 percent induced).

rules (though not across parameterizations). Mirroring the results from the unweighted cooperation rates, we again find strong support for the theoretical prediction. Total cooperation under the endogenous transition, and controlling for the differing state compositions, is significantly greater than both the static and exogenous rules ($p < 0.001$ for each comparison).

While the inference above is over unilateral cooperation rates *across* transitions regardless of the parameterization, Table 2 makes clear that *Result 1* is primarily driven by greater cooperation in the *difficult* endogenous-transition treatment. This feature of the data directly feeds into our second hypothesis, which compares cooperation-rate changes *within* transition. Looking at these within-transition shifts we summarize our conclusion from the table's data in the following summary:

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

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, the Δ column in Table 2. Our experimental results indicate a 25 percentage point reduction in cooperation (significantly different from zero with $p < 0.001$) for the static transition, and an 18 point reduction under the exogenous transition ($p = 0.009$). In contrast, endogenous transition reduction is quantitatively small (and statistically insignificant) at just 2 percentage points. Testing these drops in cooperation across transition rules we find that the reductions under the exogenous and static rules are significantly different from 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.

Looking at differences in the state-conditioned rates within transition, the relative patterns are similar, with the endogenous transition rule showing a much-smaller proportional reduction in cooperation following a move from *easy* to *difficult* than the static or exogenous rules (joint test *p*-values of 0.031 in the low state, and 0.001 in the high state). However, two additional features of the data emerge when we split the data by state: First, there *is* a small but significant reduction in high-state cooperation under the endogenous transition, though cooperative behavior is still the modal high-state response. Second, even though cooperation is theoretically possible in the low state in every treatment (and so the best-case SPE predicts no effect), we do observe significant reductions in low-state cooperation under the *difficult* versions of the static and exogenous transitions. We come back to these discrepancies in Section 3.3 when we consider learning across supergames and states. Before this though, in the next subsection we show that cooperation is driven by history-dependent play.

---

[17]Separately, the exogenous and static transition-rule treatments are also significantly different from endogenous ($p = 0.039$ and $p = 0.004$ against one-sided hypotheses).

3.2. **Cooperation Within a Supergame: History Dependence.** The evidence presented so far provides a test of our hypotheses using measures of cooperation that aggregate across the supergame. However, the mechanics of the behavioral predictions can be further scrutinized by studying choices within a supergame. Our theoretical predictions over the best-case SPE rely upon history dependence, with ongoing cooperation conditioned on past play. A first-order question is the extent that subjects respond to the history of play, where the summary result we demonstrate in this section is:

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

We start by using a reduced-form analysis tailored to the specifics of each treatment to show that behavior in each is history dependent. A shortcoming of this analysis is that it will not allow us to compare measures of history-dependent behavior across treatments. At the end of the section, we construct a measure that allows for a direct comparison across treatments.[18]

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 best-case SPE, and study how behavior in the subsequent period depends on the matched player's behavior. For the endogenous transition-rule treatments we focus on histories in which both players cooperated in period one (in the low state) and player i also cooperated in period two (in the high state as both initially cooperated). That is, we focus only on supergames with the histories $\widetilde{\mathcal{H}} = \big((C, C, L), (C, a_2^j, H) \big)$, for either period-two choice of the matched player j, $a_2^j \in \{C, D\}$.[19] 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 j cooperated or defected in period two.

The output of the regression is presented in the last column of Table 3. The likelihood of cooperation decreases in both endogenous-rule 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 *Diff-Endog* it is $0.802\ (= 0.967 - 0.165)$.[20] The reduced cooperation rates in response to the other player's defection is not statistically different across our two

---

[18]Table A5 in the Online Supplement shows the observed frequencies of the most-common histories by treatment. The main message from the 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.

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

[20]The reduction in the likelihood of cooperating in period three is significant for both endogenous-transition treatments (both $p < 0.001$)

TABLE 3. History Dependence

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

parameterizations.[21] The cooperative response in the endogenous-transition treatments is therefore predominantly history dependent.

When the transition rule is exogenous and when the state does not change through the supergame (static) we instead focus on period-two choices of subject i, conditioning on the histories where subject i cooperated in period one. That is, we use only the supergames with the histories in $\widehat{\mathcal{H}} = \left\{(C, a_1^j)\right\}$, where the first player cooperated in period one, and we will study whether the period-two choice depends on the period-one behavior of player j.[22]

In period one of exogenous-transition 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 the state realized in 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 j cooperated in the previous period, interacted with the

---

[21]As 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.
[22]This history represents 71.9 (70.2) percent of our data in the *Easy-Exog* (*Easy-Static*) treatment, and 59.1 (48.3) percent in the *Diff-Exog* (*Diff-Static*) treatment.

TABLE 4. Accuracy of Solution Concept (Last Five Supergames)

Treatment	Overall ($1 \leq t \leq 5$)			Longer-run ($t = 5$)		
	MPE (M_{DD})	Best Markov	Best SPE	MPE (M_{DD})	Best Markov	Best SPE
Easy-Endog	36.7%	64.9% (M_{DC})	81.6%	46.7%	79.0% (M_{DC})	83.8%
Diff-Endog	38.7%	61.3% (M_{CC})	74.3%	52.9%	66.2% (M_{DC})	72.9%
Easy-Exog	43.9%	57.4% (M_{CD})	86.6%	53.8%	53.8% (M_{DD})	90.0%
Diff-Exog	61.6%	61.6% (M_{DD})	76.5%	72.4%	72.4% (M_{DD})	86.2%
Easy-Static	39.6%	60.4% (M_{CC})	87.6%	47.1%	53.8% (M_{CD})	93.8%
Diff-Static	64.9%	64.9% (M_{DD})	79.1%	71.9%	71.9% (M_{DD})	86.2%
Average	47.6%	61.8%	81.0%	57.4%	66.2%	85.5%

parameterization. The results are presented in the third and fourth columns for the exogenous transition, and the first two columns for static, in each case conditioning on whether θ_2 was randomly assigned as the low or the high state.[23]

The estimates for the exogenous and static transitions show similarly strong evidence for history-dependent cooperation, at similar levels to the endogenous-transition 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.[24] 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-transition treatments if the state is low (high) the likelihood subject i cooperates increases by 57.5 (30.2) percentage points if their partner j cooperated in the previous period.[25] In low-state static supergames, the difference in the period-two cooperation rate is close to seventy percentage points as we shift the matched subject j's first-period action. At approximately eighty-percentage points, this effect is slightly higher in the high state, with similar magnitudes across the *Easy-* and *Diff-Static* treatments.

Finally, we develop a measure to place detected history dependence in context across treatments. As a first simple measure, we compute how much data in each treatment can be rationalized with the best SPE.[26] That is, we will look at the choices of each subject in a supergame and compute how closely their choices adhere to the best SPE. For example, a subject in *Easy-Endog* who cooperated

---

[23]While 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 under the endogenous transition are not the same as those under the *exogenous* and *static* transitions. The goal oinstead isto evaluate whether cooperation within each transition rule is conditioned on past play.

[24]The differences equal 0.723 ($\theta_2 = Low$) and 0.773 ($\theta_2 = High$), and each one is statistically significant ($p < 0.001$).

[25]Both figures are quantitatively large and statistically significant ($p < 0.001$).

[26]The best SPE in all *easy* parameterizations and in *Diff-Endog* is S_{DD}^{CC}. In *Diff-Exog* and and *Diff-Static* we use the on path strategy of M_{CD} which reverts to M_{DD} following *any* low-state defection, the strategy S_{DD}^{CD}.

up until an observed defection, defecting in all subsequent periods would have a 100 percent match to S_{DD}^{CC}, the best SPE in that treatment. We can then aggregate and evaluate what fraction of the data in each treatment can be rationalized with the best SPE. Because the criterion (the best SPE) is the same across treatments, we can then see if the proportion of behavior that is rationalized by the best SPE changes across treatments.

Table 4 computes the proportion of choices that are rationalized by the best SPE for two subsets of the data. The first panel uses the first five periods in each supergame and the second panel focuses only on the period-five choice (the last choice that we commonly observe in all supergames) is consistent with S_{DD}^{CC}. The table makes clear that for either measure a large proportion of the data in all treatments is consistent with the best SPE. This pattern is consistent with behavior being history dependent, as shown treatment-by-treatment in Table 3. The measure also shows that the differences across treatments are relatively small. For example, approximately eighty percent of the data is consistent with history-dependent play in the first-five periods regardless of the transition rule or parameterization. This suggests that a large majority of subjects are employing history-dependent strategy in all treatments. As Results 1 and 2 showed that cooperation rates do largely differ across treatments, this points out that the coordinative success of the history-dependent behavior does strongly depend on the transition and parameterization.

In comparison to the relative success of the best SPE, history-independent strategies such as the MPE do not rationalize the data well. The *MPE* columns show the proportion of the data that is consistent with the common history-independent equilibrium prediction, M_{DD}. Notice that a subject choosing entirely at random would have a 50 percent accuracy relative to *any* deterministic strategy, and so the average predictive power of M_{DD} is below that of random choice, suggesting poor fit. In some treatments there are better fits to other (non-equilibrium) history-independent strategies. Among the four pure-strategy Markov profiles, the "Best Markov" columns indicate the strategy with the best predictive accuracy (strategy in parentheses). Even relative to the best-fitting Markov profile the history-dependent SPE represents a significant predictive gain, increasing the predictive accuracy by approximately 20 percentage points, both for the overall data, and for the longer-run supergame outcomes. The table demonstrates that a substantial amount of the strategic heterogeneity in our games can be explained through a simple history dependent strategy. While additional gains can be made by allowing for a distribution of employed strategies, our point here is to show the additional predictive power that comes from a simple history-dependent solution concept over the non-history-dependent alternatives.

3.3. **Cross-State Effects.**

Hypotheses on Partial Cooperation. Having shown that cooperative behavior is highly 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 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 subjects' beliefs that partners will not cooperate in high-state periods reduces their 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 that continue 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.

---

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

TABLE 5. Joint Cooperation in the Low State

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

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

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

Low-State Cooperation Results. Table 5 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 period 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 insignificant and small.[29]

In contrast, for the static and exogenous transitions, 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 *Diff-Static* together with *Diff-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.

---

[28]The main findings are qualitatively the same with greater statistical confidence if we use all supergames, and if we examine unilateral cooperation.

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

In the final two columns of Table 5, 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 period, we measure the fraction that are still jointly cooperative in the low state in period five.[30] For static supergames that repeat the low-state stage-game, jointly cooperating in the first period leads to a very high likelihood of being jointly cooperative in period five. Approximately 90 (80) percent of the *easy* (*difficult*) repeated games are persistently cooperative, and the difference between parameterizations is not significant. In contrast, while the *Easy-Exog* game has 80 percent of the initially cooperative supergames continuing to jointly cooperate in period five, this falls to a little over 40 percent for the *Diff-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 for the static-rule low-state game across parameterizations. Table 5 makes clear that there is a significant reduction in initial low-state joint cooperation with the static transition rule, with joint-cooperation rates of 55 and 32 percent in *Easy-Static* and *Diff-Static*, respectively.[31] But conditional on a low-state being initially selected the supergames are structurally identical, as the parameterization change only affects the high-state payoffs. We now show that this is a learned response; that the reduced cooperation in the high-state games has a contagious effect across supergames.

To start with, in Table 6 we present data on the very first decision subjects make. The table provides the unilateral cooperation rates by treatment in the low state, with standard errors and inference from a regression of a treatment dummy for cooperation on treatment dummies. We constrain our attention to static treatments first. Notice that while Table 5 shows more low-state cooperation towards the end of each session in *Easy-Static* than *for Diff-Static*, there is essentially no difference at the beginning of the session. Per Table 6, the unilateral cooperation rate in the low-state for the first static-rule supergame is actually higher in the difficult parameterization—64 percent compared to 50 percent in *easy*, though the difference is insignificant. We now show that differential response at the end of sessions is driven by subjects' experiencing the high-state supergame.

---

[30]On period five we condition on the subject pair successfully cooperating in period one, and the game being in the low state. We focus on period-five behavior because this is the last period of the block design for which we have data for all supergames. We do not present results for the endogenous transition as being in the low state is not independent of behavior, necessitating a joint defection 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 period five. This compares to 32 percent of *Diff-Endog* supergames, with the reduction highly significant ($p = 0.006$).

[31]Initial unilateral cooperation rates are provided in Table A2 in the Online Supplement.

TABLE 6. Initial Unilateral Cooperation (First Supergame, $t = 1$, $\theta =$Low)

Transition	Easy	Diff.
Static	0.500	0.643
	(0.086)	(0.126)
Exogenous	0.571	0.381
	(0.073)	(0.073)
Endogenous	0.762	0.762
	(0.073)	(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 (across the first ten session supergames) we can identify the cross-state effects. As a simple specification, we examine the decision to initially cooperate in low-state supergames at the end of the session (the last five supergames). We regress these unidecisions 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]

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

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-level environments in isolation. Instead the selection of cooperative outcomes responds to the session-level environment.[34]

The information in Table 6 also shows that there is a difference between endogenous and the other transition rules, even from the beginning of the session. The endogenous-transition treatments

---

[32]Theoretically, 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:

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

[33]Unlike 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.

[34]Looking 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.

have significantly greater cooperation from the very first subject 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 through the instruction's description of the environment.

We summarize the findings as:

Result 5: While subjects do internalize the potential gains from an endogenous transition early on, they do not otherwise display significant differences in initial low-state cooperation. Across the session however, there are significant responses across environments, and to the strategic dynamics within environment. In static transition-rule 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 the future environment. 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 of agent's actions.

While our results show reduced cooperation in the high state as we manipulate the strategic tensions within it, 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 the treatment does not support cooperation in every state our results indicate substantial cooperative reductions in all states. The observed contagions are found both in initial play, and in the ability to sustain joint cooperation in the long run. While our results are mostly optimistic on subjects' ability to internalize

---

[35]A joint test finds that the endogenous-transition treatments have significantly greater cooperation than in either the static- or exogenous-transition 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* treatments yields a failure to reject equality ($p = 0.237$).

strategic 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 co-operation can be supported in the worst-case state. Rather than the partially cooperative outcomes that can be supported by history-dependent play, 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.

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

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

Bó, 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*, 2018.

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.