



Choosing a future based on the past: Institutions, behavior, and path dependence



Jenna Bednar^{a,*}, Andrea Jones-Rooy^b, Scott E. Page^{c,**}

^a University of Michigan, Ann Arbor and Santa Fe Institute, 426 Thompson Street, Ann Arbor, MI 48104, USA

^b New York University Shanghai, 1555 Century Avenue, Pudong, Shanghai 200122, China

^c University of Michigan, Ann Arbor and Santa Fe Institute, 1085 S. University Ave., Ann Arbor, MI 48109, USA

ARTICLE INFO

Article history:

Received 12 September 2014

Received in revised form 15 September 2015

Accepted 28 September 2015

Available online 8 October 2015

JEL classification:

C630

D830

Keywords:

Culture

Behavioral spillovers

Institutional sequencing

Development

Agent-based models

Computational modeling

ABSTRACT

Institutions do not always produce behavior consistent with what theory predicts, leading comparative scholars to turn to explanations based on historical or cultural exceptionalism. Context can influence not only how an institution performs but also the very choices of institutions that societies choose to govern themselves. In this paper, we construct a model that produces contextual effects that result in *institutional path dependence*. In doing so, we provide formal foundations for qualitative arguments that context matters and identify a contributing causal mechanism: *behavioral spillovers*. Using both mathematical and computational techniques, we show that spillovers provide a mechanistic explanation for how pre-existing institutions affect the performance of new institutions as well as the optimal choice among institutions. We find that these spillovers can depend on either the set or the path of previous institutions. Both results support qualitative arguments that historical institutional contexts influence outcomes in current institutions. Importantly, the spillovers depend not only on the outcomes produced in the institutions but also on the specific behavior that produces the outcomes. As a result, we show that institutions that create diverse ensembles of behaviors generate better outcomes and less path dependence than those that cause all agents to converge on the identical strategy.

© 2015 Published by Elsevier B.V.

1. Introduction

Individuals interact in multiple contexts ranging from formal political and economic institutions to informal organizational and social settings. Looking across societies, *Ostrom noted what she called the “diversity of regularized social behavior” (2005:6)* as well as the diversity of institutional forms and structures used to allocate resources and opportunities. Many of these institutions induce behaviors that result in efficient, socially desirable outcomes, but many also fail. In some cases, the failure results from poor design. Some institutions endow individuals with misaligned incentives or insufficient information to take optimal actions.¹

In other cases, even though the institution creates proper incentives and appears to endow individuals with enough information, behavior doesn't align with what was expected. These violations occur both in controlled experiments and in practice. This apparently suboptimal behavior presents a puzzle. As *Smith (2003)* notes in his Nobel Lecture: “If people in certain contexts make choices that contradict our formal theory of rationality, rather than conclude that they are irrational, some ask why, reexamine maintained

* Corresponding author. Tel.: +1 734 615 5165; fax: +1 734 647 4575.

** Corresponding author. Tel.: +1 734 647 9193; fax: +1 734 763 9267.

E-mail addresses: jbednar@umich.edu (J. Bednar), spage@umich.edu (S.E. Page).

¹ In the language of mechanism design, such institutions fail to implement the social choice correspondence (*Reiter, 1977*).

hypotheses including all aspects of the experiments—procedures, payoffs, context, instructions, etc.—and inquire as to what new concepts and experimental designs can help us to better understand the behavior” (2003:471).

Attempts to understand apparently suboptimal behavior in institutional settings generally take one of two approaches. Some, including Smith, look in more detail at the informational and computational requirements of the mechanism: do people know what they need to know and can they compute the equilibria? Others abandon the rational choice paradigm altogether in favor of a behavioral economic approach (Akerlof and Shiller, 2009; Camerer, 2003). Both approaches can point to successes but each has shortcomings. The first approach obliges a case by case analysis. The second approach promises a more general theory, one that replaces *homo economicus* with actors whose behavior is thought more psychologically plausible, but that behavior may be computationally more taxing.

Neither approach has much to say about cultural variations in behavior found in bargaining games (Roth et al., 1991) and the Ultimatum Game (Henrich et al., 2004, 2010). Some cultural groups exhibit more prosocial behavior than others and that behavior does not appear to be solely the product of an innate psychology, but a reflection of norms and institutions that have emerged over the course of distinct cultural histories. Both behavior and institutional performance can depend on context (Nisbett et al., 2001).² To test for the existence and magnitude of contextual effects, experimental social scientists have begun to study how agents react when they play multiple games (Bednar et al., 2012a,b; Cason et al., 2012).

In this paper, we construct a model in which ensembles of games are built up over time. The sequential construction of an ensemble creates a history of behaviors and experiences that may spill over into subsequent strategic contexts. To quote Dolan and Galizzi (2015), “no behavior sits in a vacuum.” Context matters. In this paper, we explore theoretically whether a specific type of spillovers across games can produce behavioral diversity within common games in distinct ensembles.

We also explore whether these spillovers can affect the choice of institutions. Agents often have a choice over institutional structures. In many settings, ensembles of institutions are built up over time. We therefore also investigate whether the behavior spillovers influence, or bias, that choice over institutions. These influences may produce either set or path dependence—that is, they may make institutional choices dependent on the other elements of the ensemble (set dependence), or on the order the other institutions were introduced (path dependence).

To capture these behavioral spillovers, we assume that individuals' initial behaviors depend on past behaviors, but that individuals then learn. Ultimate behavior will be rational in so far that the individuals are unable to learn a behavior that leads to higher payoffs. The assumption that when confronted with a new situation, people may, at least initially, behave similarly to how they have behaved in a related context has been analyzed in the context of decision theory (Gilboa and Schmeidler, 1995).³ Experimentalists also find evidence of spillovers of this sort. To again quote Smith, commenting on decades of research: “the challenge of any unfamiliar action or problem appears first to trigger a search by the brain to bring to the conscious mind what one knows that is related to the decision context” (2003:469).

In our model, the spillovers only affect initial actions. Final outcomes depend on an interaction between the learning rule and the initial action distribution.⁴ In our analysis, we show mathematically that learning can produce efficient equilibria in the games that we study but that it need not (Camerer, 2003). Our interest, and the focus of our analysis, is on how the initial actions—which are a product of past games—influence outcomes.

One of our main findings will be that how individuals play one game can depend on what games they played previously. Thus, our model provides a candidate explanation for behavioral variations in common institutional settings. That explanation is based on spillovers that result from existing behavioral repertoires.⁵ Some societies may be predisposed to perform poorly in some institutional settings because their initial behaviors lead them to inefficient equilibria. It follows that optimal institutional choices can also depend on the sequence of past institutions. In other words, behavioral spillovers can produce set and path dependent choices over institutions.

As mentioned, our framework has its origins in mechanism design which applies the canonical game-theoretic approach to modeling institutions (Hurwicz, 1972; Hurwicz and Schmeidler, 1978). Mechanism design characterizes institutions as action sets, payoffs, and communication structures. Optimal institutions align individual incentives with the collective interests and aggregate privately held information to the extent possible given constraints on participation and misrepresentation. In brief, institutions create incentives to advantage some behaviors over others and for the revelation of information that is of collective value or interest. Mechanisms can then be evaluated by their ability to produce good outcomes in equilibrium (see also Diermeier and Krehbiel, 2003). Here, we look not just at the outcomes but also at the behavior that produces those outcomes.⁶

Our framework extends the standard mechanism design approach in three directions. First, we consider *multiple institutions in sequence*. Second, rather than rely on equilibrium refinements, such as efficiency, symmetry, or strategic dominance to select among equilibria, we use *learning to select among equilibria*. Third, we introduce behavioral spillovers by assuming that *initial behaviors depend on their actions in similar games*.

² Evidence of significant cultural differences combined with the fact that many of the documented psychological biases have been shown to exist mostly in Western societies have led to calls for empirical tests of whether psychological biases are generic or hold only in Western countries (Medin et al., 2010).

³ In the United States, the number of offspring in immigrant families lies between the average number in their home culture and the U.S. average (Fernández and Fogli, 2006).

⁴ See Fudenberg and Levine (1998), Camerer (2003), Morton and Williams (2010), Bendor et al. (2003), and Golman and Page (2010) for various studies of how outcomes depend on the learning rule and the initial state of the system.

⁵ See also Bednar et al. (2010) for a model of where individuals seek consistency across games.

⁶ See Page (2012) for a survey and critique of mechanism design along the lines implicit in this paper.

Our conceptualization of spillovers is a novel contribution. With spillovers between games, we are able to assume that individuals do not confront a new institution with a blank slate. Spillovers can operate in two ways. First, an initial behavior might be an equilibrium behavior in the new institution, and that equilibrium need not be efficient. In this case, the spillover plays the role of an equilibrium selection device.⁷ Second, an existing behavior may fail to be an equilibrium behavior in the new institution, but can affect which equilibrium is chosen. We will in fact find that initial predispositions can have large and unexpected effects. The cumulative effect will be the potential for set and path dependence at three levels: in behavior, in outcomes, and in optimal institutional choices.

We have three main results. First, we find strong evidence that both the *set* of existing institutions as well as can the *order* in which they arise can influence behaviors. Thus, the model produces behavioral diversity in common games based on institutional context. Some of our findings will be intuitive: if agents learn to cooperate, be selfish, or alternate in one game then those behaviors will be more likely to occur in subsequent settings. However, not all spillovers prove to be so straightforward. We find that learning to alternate can lead to more selfish behavior in a Stag Hunt game than will learning to be selfish. These counterintuitive results can be understood by looking at behaviors and not just outcomes. For example, how individuals cooperate in one setting (using Tit for Tat or Grim Trigger) can influence their capacity to cooperate in another setting.

Second, we find that the optimal institutional choice for a community can depend on the *set* of previous institutions as well on the *order* the institutions were put in place. The best institutional choice cannot be made without regard to existing institutions. While many models of path dependence rely on increasing returns or positive feedback (Pierson, 2000, 2004), our model does not require that assumption, in keeping with a more expansive theory of path dependence laid out in Page (2006). Third, some early institutional choices allow more flexibility than others over future institutional choices. We formalize this idea by introducing the concept of *path limiting behaviors*. These are behaviors that prevent efficient learning in some future games, i.e. they limit the path of what is possible.

In our analysis, we hold constant any differences in how members of a community learn, a feature that might also affect institutional performance. One community may be more individualistic, another may be more collectively-minded (Inglehart, 1977). It has been shown that individual learning rules and collective learning rules differ in the equilibria they locate (Vriend, 2000; Golman and Page, 2010). Thus, societies that learn differently may experience distinct outcomes when using identical institutions to allocate resources.

We confine our analysis to the model itself with the aim of demonstrating that behavioral spillovers can produce behavioral diversity in common games as well as set and path dependent choices over institutions. Obviously, these findings have implications for area studies generally as well as the substantial literature on institutional failure. In the discussion at the end of the paper, we interpret our findings more expansively within these two contexts.

2. A model of sequential institutions and behavioral spillovers

In our model agents confront a sequence of institutions that we formalize as games. Each game is played repeatedly. The agents first learn to play the initial game in this sequence. They then add a second game, and eventually a third. The spillovers arise because the agents' initial strategies in the second game depend to a varying extent on the strategies that they learned in the first game. We refer to the extent of that dependence as the *behavioral spillover*. In the most extreme case, every agent initially plays the new game using the strategy it learned in the first game. The agents then learn how to play in this second game. This construction produces a *vertical ensemble* of games in which the learned behaviors in the second game can depend on the strategies used in the first game. This differs from a situation in which the agents confront the entire ensemble of games at the same time, or a *horizontal context* (Bednar and Page, 2007).⁸

This vertical influence on behavior and therefore outcomes creates the possibility of both *set* and *path dependent* behaviors. Set dependence holds if the learned behavior in later games depends on the games previously introduced. Path dependence is a stronger condition, arising if the same two games played in the opposite order produce different behaviors. That is, with path dependence, order matters. If either set or path dependence arises, then it also follows that the optimal choice of a game form could be a function of the set or path of previous choices. Therefore, in addition to *behavioral dependence* (the notion that behaviors depend on the set or order of earlier games), there may also exist what we call *institutional dependence*: the optimal choice of an institution, i.e. the game form, may depend on the set or order of previous games.

It is of course possible that the learned behavior in the second game does not depend in any way on the strategies developed in the first game. If so, then we will say that the second game is *immune* to behavioral spillovers. We will find that games with dominant, Pareto efficient strategies are immune, but that games that have multiple equilibria in the repeated game setting tend not to be immune.

2.1. Set and path dependence

To lay the groundwork for our formal model, we present examples of set and path dependence. These examples demonstrate how behavior can bleed from one game to the next. We begin with an example of set dependence. In this example, the second game will be what we call the *Knife Edge game*.

⁷ Our emphasis on behaviors differs from the culture-as-equilibrium selector literature, where beliefs, not behavior, are the mechanism that selects the equilibrium. See, for example, Greif, 2006.

⁸ Ensembles of games are collections of independent games played by the same set of agents; agents may apply the same strategy across games but they may also choose to apply different strategies. In this sense, *ensembles* differ from nested games in which one action has implications in many games (Tsebelis, 1990).

Knife Edge game			
		Column	
		C	S
Row	C	8,8	2,14
	S	14,2	4,4

The repeated version of the Knife Edge game has many equilibria. Here and throughout the paper, we consider equilibria of the infinitely repeated game with a discount rate near one. In our computational experiments, we run the games for a finite number of rounds. Given that the computer agents do not know the number of rounds, it is as if they are playing an infinitely repeated game. For the purposes of this example, we focus on two equilibria: the *cooperative equilibrium*, in which both agents choose C in each round, and the *alternating equilibrium*, in which the agents alternate between the outcomes (S, C) and (C, S). In each of the first two equilibria, each agent receives an average payoff of eight.⁹

Recall that the Knife Edge game is played second. We want to see if behavior in that game depends on the first game. Let's suppose that the agents first play the *Prisoners' Dilemma*. In this game, the agents may well learn to cooperate.

Prisoners' Dilemma game			
		Column	
		C	S
Row	C	8,8	1,10
	S	10,1	2,2

Assume a maximal level of the behavioral spillover, so that when the Knife Edge game is added, the agents initially play the cooperative strategy that they were using in the Prisoners' Dilemma. This might be Tit for Tat, or it might be Grim Trigger. The agents then learn. The strategy used in the Prisoners' Dilemma might become entrenched in the second game. The agents might learn to play the cooperative equilibrium in the Knife Edge game as well. If so, we will say that the learned behavior in the first game spills over into the second game.

Suppose that instead, the first game played was what we call the *Alternation Game*. In this game, the efficient equilibrium strategy calls for each pair of agents to alternate between (C, S) and (S, C). We refer to this as the *alternation strategy*. If the agents initially play the alternation strategy in Knife Edge, then even after learning, they may stick with this strategy.

Alternation game			
		Column	
		C	S
Row	C	2,2	-2,10
	S	10,-2	2,2

The table below shows how learned behavior in the initial game can influence the learned behavior in the second game through the spillover.

Ensemble composition affects behavior		
	Community 1	Community 2
Game 1 (Behavior)	Prisoners' Dilemma (Cooperate)	Alternation (Alternate)
Game 2 (Behavior)	Knife Edge (Cooperate)	Knife Edge (Alternate)

⁹ Note that we assume no discounting.

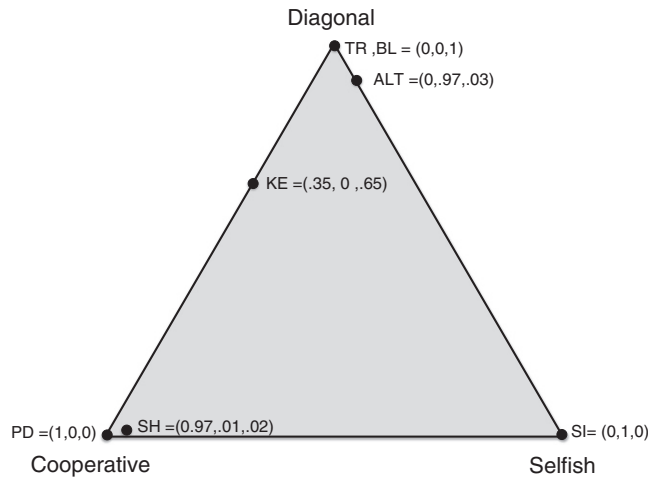


Fig. 1. Simplex representation of the outcome distributions for the seven games (400 trials).

In this example, behavior differs in the Knife Edge game across the two communities because the initial games differ and those initial games produce behaviors that then get applied in the Knife Edge game. As this example suggests, we are using the learning rules as a type of refinement criterion. By the Folk Theorem, any repeated game has lots of equilibria. We're using an evolutionary learning rule to select from among those equilibria. The initial conditions for the rules are influenced by the existing set of games.

To show *path dependence*, we must show that the behaviors depend on the order in which the games are played. To see how this could occur, assume that in one sequence the Knife Edge game is played first and the Alternation Game is then added. Suppose that in both games, individuals learn to play cooperatively, i.e. to take the action pair (C, C). Assume next that we switch the order so that the Alternation Game is played first. Now, suppose that the individuals learn to alternate in the first game and then use the same strategy in the second game. Thus, the same games, in different orders, produce different behaviors.

Path dependent behavior		
	Sequence 1	Sequence 2
Game 1 Behavior	Knife Edge (Cooperate)	Alternation (Alternate)
Game 2 Behavior	Alternation (Cooperate)	Knife Edge (Alternate)

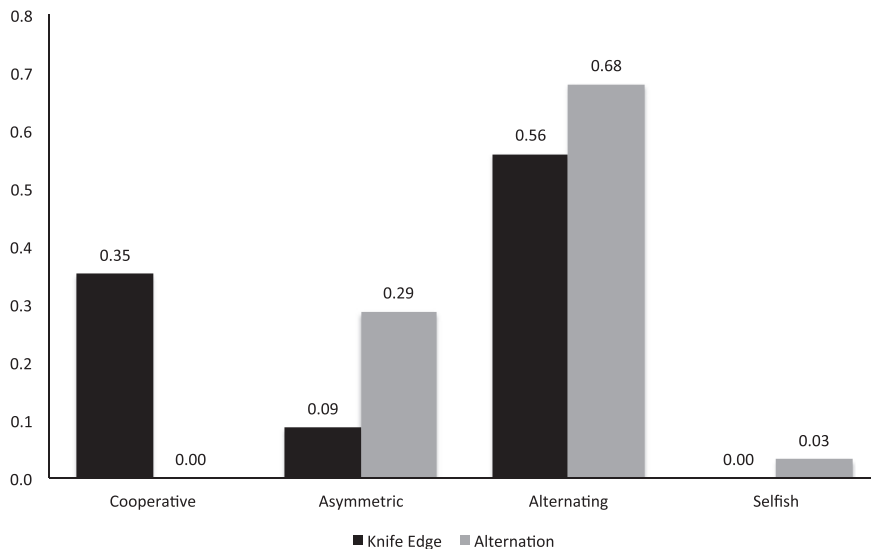


Fig. 2. Strategies used in Knife Edge and Alternation game (400 trials).

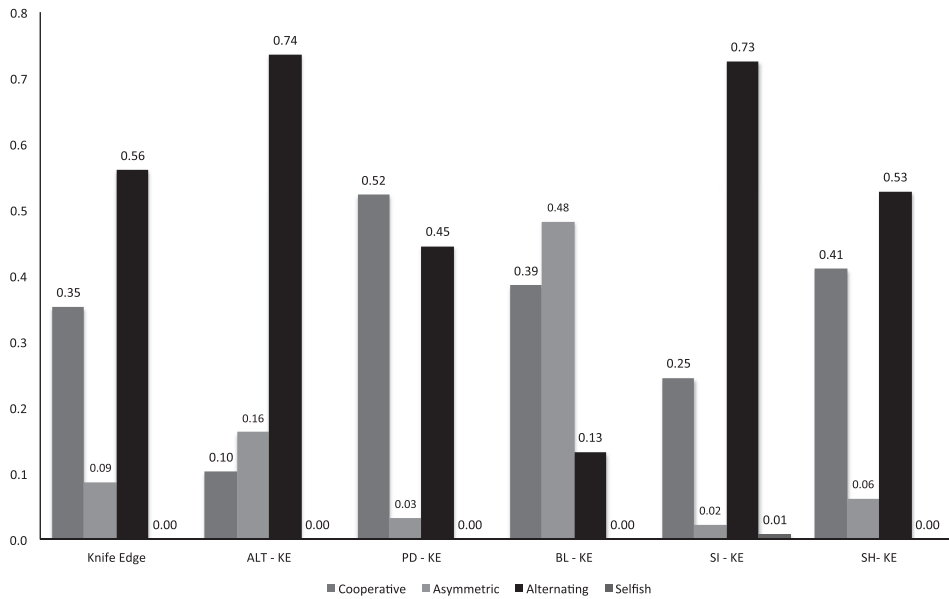


Fig. 3. Outcomes in Knife Edge under different initial institutions (400 trials).

Although we have only considered two games so far, we could add a third game. When adding a third game, we must make some sort of assumption about how initial behavior depends on the past behavior. If the agents use the same strategy in each of the first two games, then we can assume that they use that strategy initially in the third game. If the strategies differ, then we can assume that the agents are more likely to use the strategy of the game that's more similar to the new game. In this case, that would be the Knife Edge game.

If we assume that all agents cooperate in both games in the first sequence, they we might well expect cooperation in the Prisoners' Dilemma. If the players alternate in each of the first two games, then it's less clear what the agents will learn to play in the Prisoners' Dilemma. They could alternate. They could cooperate. Or, they could defect.

2.2. The model

We now present our formal model. We assume a finite population of agents who interact in multiple institutional contexts that we model as repeated games. New institutions are added sequentially. Agents first play one game, and then a second game is added. Agents learn to play each new game until another game is added. Therefore, each era in which a new game is introduced consists of numerous periods each containing multiple rounds of play giving agents the opportunity to learn effective behaviors in the new institution before another institution is added.

Within this general framework, the game forms could include any number of actions or strategies and involve any number of players. Here, we restrict attention to seven two-player game forms. Each game has two actions, one that is more selfish (S) and one that is more cooperative (C). In this way, it makes sense to think of agents transferring behaviors from one game to another.

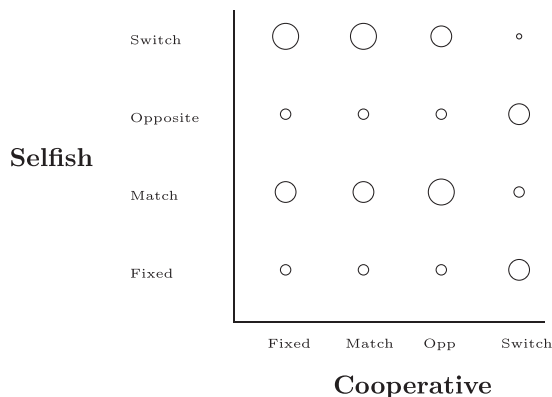


Fig. 4. A randomly drawn population of initial behaviors.

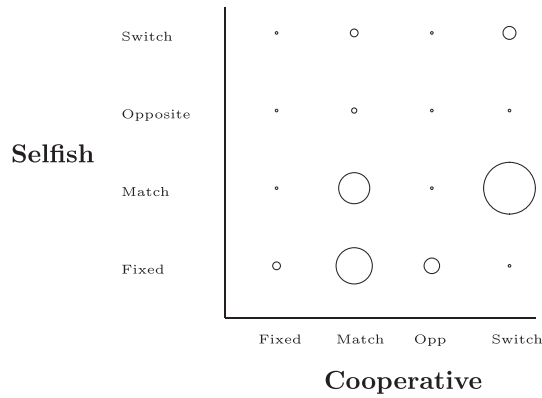


Fig. 5. Behavior: Knife Edge.

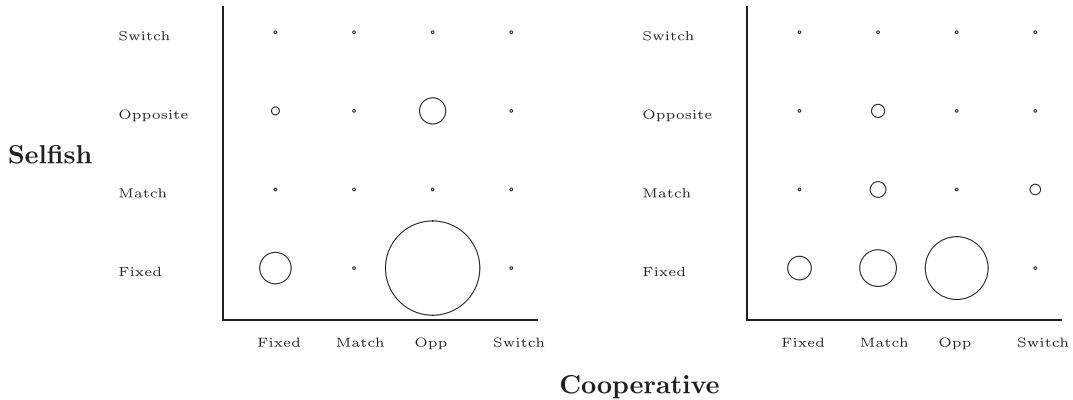


Fig. 6. Behavior: Bottom Left and Knife Edge following Bottom Left.

We have already described three of the games: the *Knife Edge*, *Alternation* and *Prisoners' Dilemma* games. The fourth game is the familiar Stag Hunt game in which the selfish strategy has a guaranteed payoff, but cooperation gives a higher payoff only if the other player cooperates as well. We call the fifth game *Self Interest*. In this game, choosing to be selfish is a strictly dominant strategy.

Self Interest (SI)				Stag Hunt (SH)			
		Column				Column	
		C	S			C	S
Row	C	0,0	0,6	Row	C	8,8	0,6
	S	6,0	8,8		S	6,0	6,6

The final two games are asymmetric, advantaging either the row or column player. Each also has two pure strategy equilibria in which one player cooperates and the other plays selfishly. We refer to these as *Top Right* and *Bottom Left*. The games represent institutions in which particular roles, say having higher status or being a manager are endowed with advantages. By including these games in the analysis we can explore whether behaviors that advantage one player learned in one institution spread to later games that have symmetric payoffs.¹⁰

Top Right game (TR)				Bottom Left game (BL)			
		Column				Column	
		C	S			C	S
Row	C	4,6	6,10	Row	C	6,4	2,8
	S	8,2	4,6		S	10,6	6,4

¹⁰ All of these games with the exception of Stag Hunt are analyzed in Bednar and Page (2007).

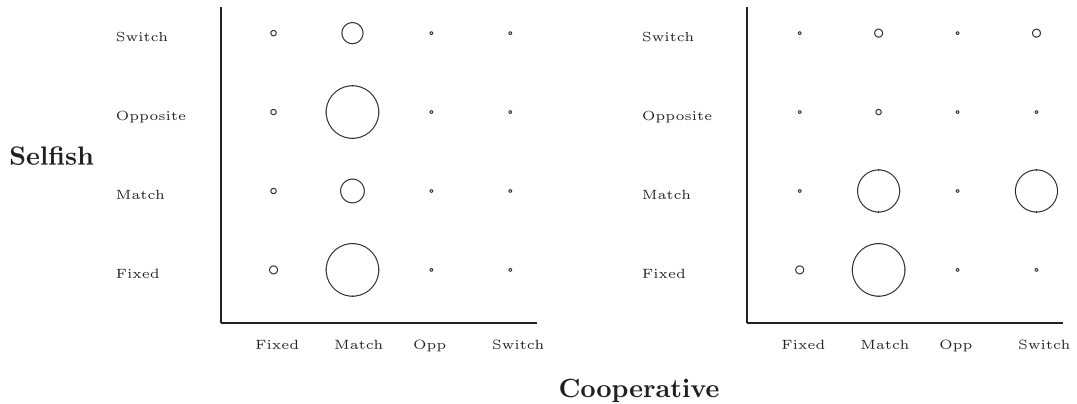


Fig. 7. Behavior: Prisoners' Dilemma and Knife Edge following Prisoners' Dilemma.

Note that all seven games have a maximal joint payoff of sixteen. We will denote the *efficiency* of the outcomes at the end of the learning process as the average total payoff divided by sixteen.

We capture the strategies of the agents using finite state automata that can encode common *repeated game* strategies such as tit for tat, grim trigger, alternate, and always defect. In games such as those we consider here, they have proven capable of generating outcomes that qualitatively align with what people produce in laboratory experiments (Rubinstein, 1986; Kalai and Stanford, 1988). The use of the automaton allows us to effectively take snapshots of the agent's behavioral rules.

Our specific formulation assumes that for each game, each agent uses automata that consist of six binary variables known as bits. When playing a game, an agent is designated as either the row player or the column player. Two of the bits denote the agent's initial action (either C or S) depending on the agent's position, either row player or column player. The other four bits denote the agent's action as a function of its action and its opponent's action in the previous round. There are four such possible pairs: CC, CS, SC, SS.

The strategy of agent *i* in game *G* $\theta_i(G)$ can be written as follows:

$$\{\text{Row 1}^{\text{st}} \text{ Action}, \text{Column 1}^{\text{st}} \text{ Action}, \text{Action}(CC), \text{Action}(CS), \text{Action}(SC), \text{Action}(SS)\}$$

where *Action*(CC) denotes the action the agent would take following the outcome CC.

To provide some intuition for this representation of strategies, the later two pairs of bits can also be thought of as capturing behavior in two *mental states*: a *cooperative* state and a *selfish* state. Given how we represent strategies, the third and fourth bits describe how the agent will behave in the next round following a cooperative action by the other agent. If the third and fourth bits are set to C, this means that once the agent cooperates, it will cooperate forever. The fifth and sixth bits represent the selfish mental state. If the agent's fifth bit equals C and its sixth bit equals S, then following a selfish action, the agent will copy the behavior of its opponent. Later, we will characterize this particular behavior as *matching*. Note that there exist sixty-four unique automata and that each encodes a reasonable strategy. For example, Tit for Tat is represented as {C, C, C, S, C, S}, Grim Trigger is represented as {C, C, C, S, S, S}, and Win Stay, Lose Shift, is represented as {C, C, C, S, S, C}.

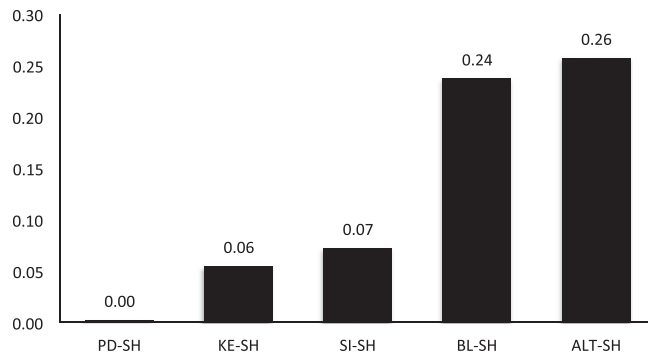


Fig. 8. Percentage of selfish outcomes in Stag Hunt (400 trials).

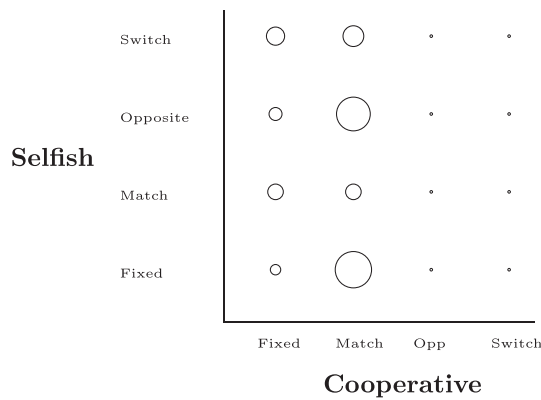


Fig. 9. Behavior: Following Stag Hunt.

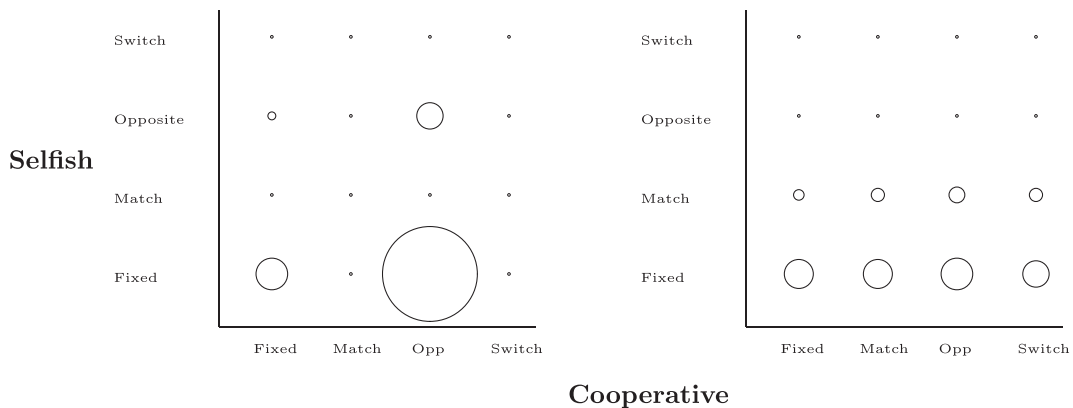


Fig. 10. Behavior: Bottom Left and Self Interest.

2.3. Evolving strategies

We assume a population of sixty-four agents. We arrange the agents in a circle and assume that each agent plays both as row and column player against the two agents to its left and the two agents to its right. Having agents play in a network increases the likelihood that cooperative and alternating strategies will emerge.

To evolve strategies we introduce two mechanisms: *mimicry* and *mutation*. With mimicry, an agent compares its total payoff to those of its four opponents, and copies the automata with the highest payoff. Mutation, by contrast, is like the random experimentation of the natural world. With a low probability (2%), each bit in each agent’s automaton is randomly reassigned.

Agents play for 150 periods. Each period is a set of four interactions—one with each near neighbor—lasting forty rounds. Following each period, agents mimic the highest-payoff automata from among its neighbors. In the first 100 periods agents also mutate. In the final 50 periods, agents do not mutate to give the population an opportunity to converge on the common best strategies.

3. Results

Our results combine mathematical theorems that establish general properties of our framework with computational experiments that show the likelihood of the various theoretically possible results. The computational model complements the mathematics in three ways. First, the mathematical theorems characterize properties of the outcomes that could occur, but as they cannot predict the likelihood of different equilibria, the simulations provide a sense of what is likely to occur. Second, the computational model enables us to produce graphs and charts that explicate the logic of the model and to explore rates of convergence and stability. Equilibrium models show existence of equilibria. Computational models show attainability and stability (Miller and Page, 2007).¹¹ And finally, the computational models let us explore the space of possible sequences of games much more widely than if we tried to prove theorems for each case.

¹¹ One can show attainability and stability of equilibria mathematically but often a new proof is required for each behavioral rule. Most papers on institutions as games show neither stability nor attainability.

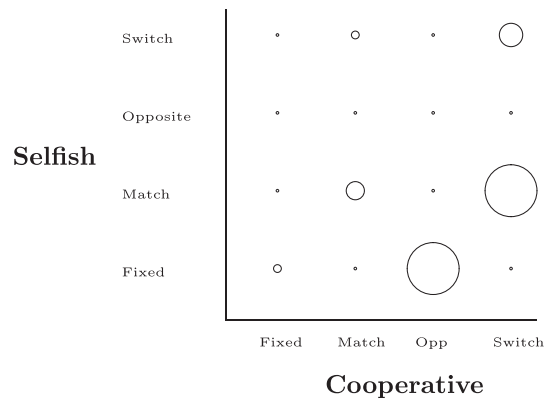


Fig. 11. Behavior: Following Alternation Game.

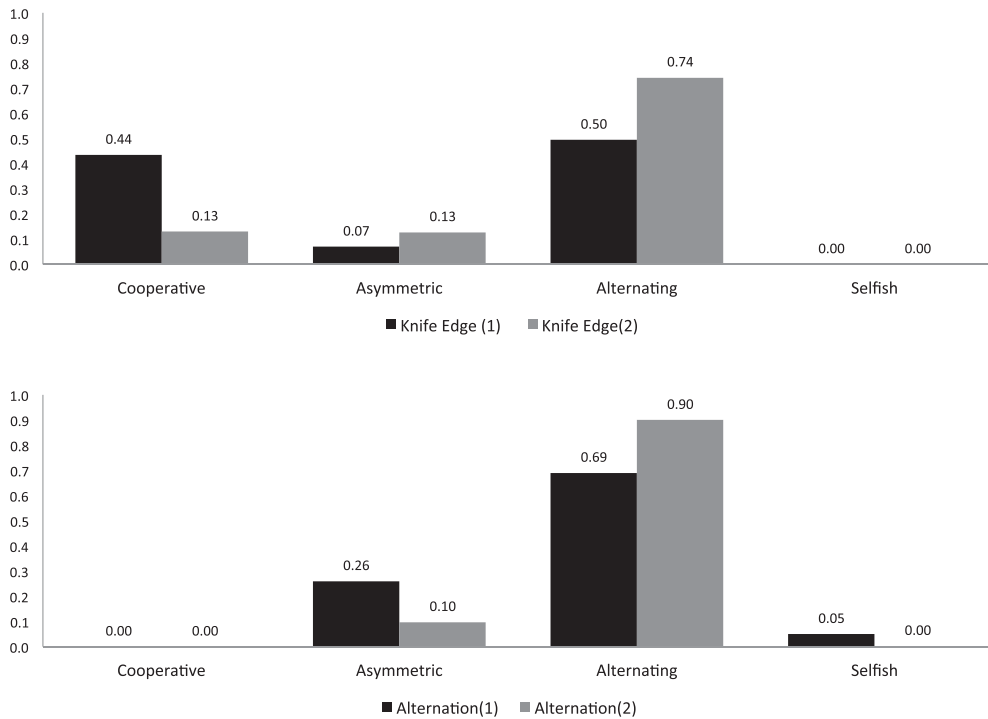


Fig. 12. Path dependent outcomes in Knife Edge and Alternation (400 trials).

3.1. Foundational mathematical results

We first state two formal claims that demonstrate that our framework can produce efficient as well as inefficient outcomes. Recall that our agents' strategies are represented as automata. During the learning process, agents both mimic the automata of neighbors with higher payoffs and mutate their automata, which are meant to capture random experimentation. In our framework, the relevant notion of an equilibrium is an *Automata Nash Equilibrium*, a set of automata, one for each agent, such that no agent could earn a strictly higher payoff using a different automaton. Given our construction, for most games there exist multiple automata that produce the same outcome so that there will not generally be strict Nash Equilibria. For example, if the entire population plays *Tit for Tat* in the Prisoners' Dilemma, an agent could switch to *Always Cooperate* and receive the same payoff.¹²

¹² As will be clear when we present our computational findings, given that we allow for mutation, our model produces *neutral drift* in behaviors and we see substantial heterogeneity in automata. For example, in the Prisoners' Dilemma a region of cooperators may include some who play *Grim Trigger*, some who play *Always Cooperate* and some who play *Tit for Tat*. An *Always Cooperate* agent surrounded by agents playing *Tit for Tat* will earn an efficient payoff.

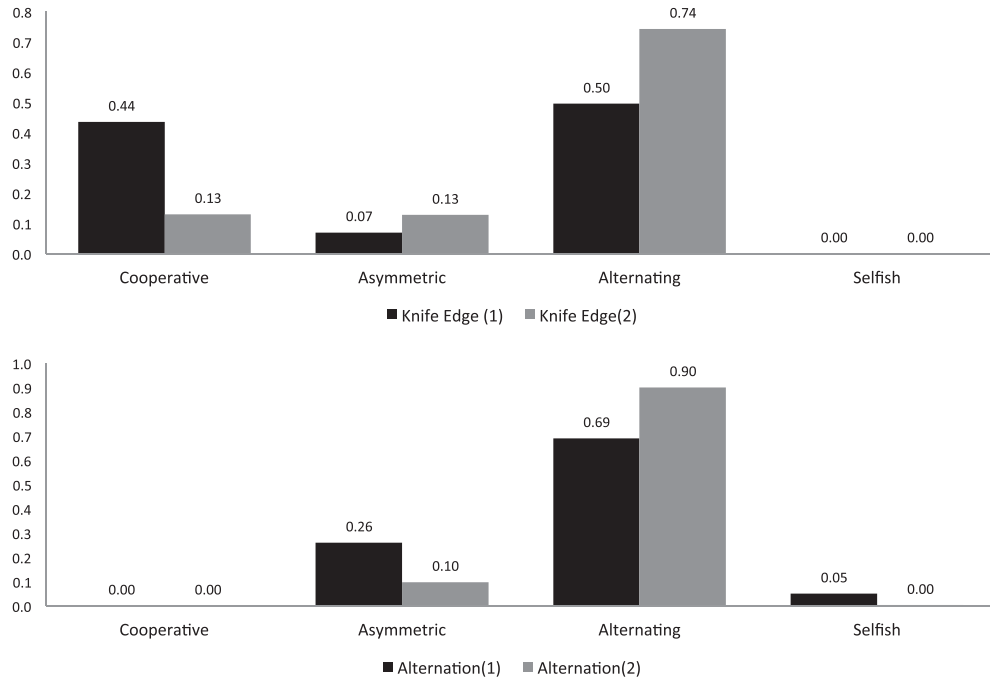


Fig. 13. Path dependent outcomes in Stag Hunt and Alternation (400 trials).

We make the following assumption which is consistent with our computational implementation of the model.

Circular Play. The agents are arranged in a circle and in each period play the game for M rounds against the two nearest agents in each direction.

Our first theoretical result demonstrates that for each game there exists an efficient Automata Nash Equilibrium. The proof along with all subsequent proofs is in Appendix 1.

Claim 1. Given Circular Play, for each game there exists an efficient Automata Nash Equilibrium.

This first claim states that it is possible that our computational experiments could produce efficient outcomes in each game. Our second claim states that it is possible that our computational experiments produce populations in which some agents' automata play efficient equilibria and some do not. Recall that in our computational experiments, there exist two stages of adaptation. In the

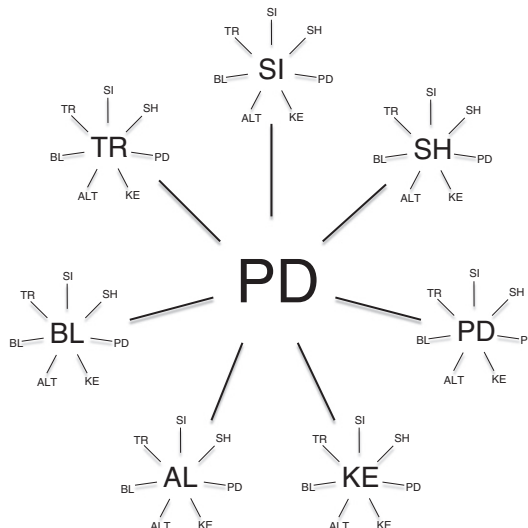


Fig. 14. The forty-nine possible paths following Prisoners' Dilemma as the initial game.

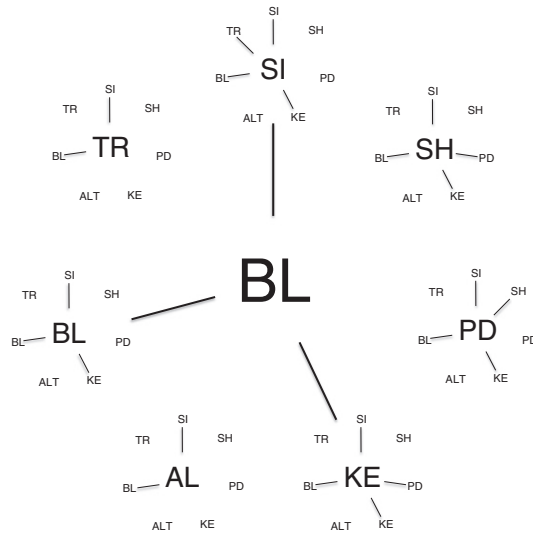


Fig. 15. The efficient paths following Bottom Left as the initial game (400 trials).

first stage (the first one hundred periods), the agents both mimic better strategies and mutate to new strategies. In the second stage (the final fifty periods), they can only mimic the automata of neighbors.

Given this construction, the automata produced by the model need not be Automata Nash Equilibria; they need only be what we call *Copy Best Neighbor Equilibrium* in which each agent either earns the largest payoff of any of the agents with whom it plays or if its automata is the same as that of one its neighbors who earns the highest payoff. In a Copy Best Neighbor Equilibrium each agent uses the automata that produce the highest payoff for any agent in its neighborhood.

Copy Best Neighbor Equilibria need not be efficient. In fact, all agents using the same automata are by definition a Copy Best Neighbor Equilibrium. Our next claim states that for each game that we consider, there exists a Copy Best Neighbor Equilibrium in which some agents produce efficient outcomes and some do not. This means that it is possible that during the first stage of learning, some agents could discover and play efficient strategies, yet these need not take over the entire population. More formally stated, there exist Copy Best Neighbor Equilibria that include both efficient and inefficient play.

Claim 2. *Given Circular Play and at least ten agents, in each game there exists a Copy Best Neighbor Equilibrium in which only a fraction of agents produce efficient outcomes in each round.*

These two results together imply that our learning model could produce efficient outcomes in each but it's also true that even if it locates automata capable of producing those outcomes, they need not take over in the population.

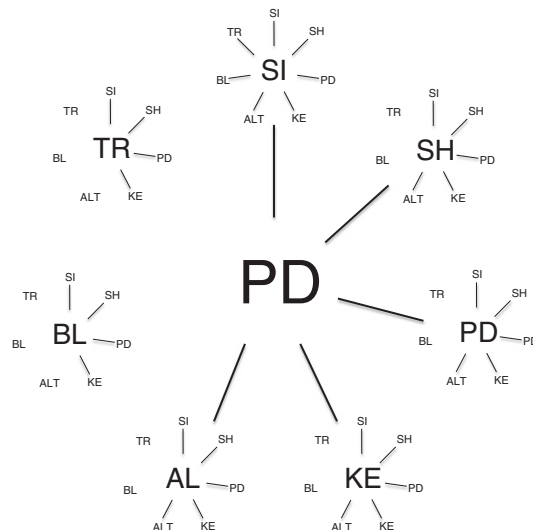


Fig. 16. The efficient paths following Prisoners' Dilemma as the initial game (400 trials).

3.2. Computational results

We first present outcomes for the seven individual games assuming a uniform initial distribution over strategies. These results provide a baseline from which we can identify set and path dependent outcomes and behaviors. We classify outcomes as *cooperative*, *selfish*, *diagonal: alternating*, or *diagonal: asymmetric* using the following procedure. We first sum the percentage of outcomes in which the two agents took different actions. We refer to these as *diagonal* outcomes. We then compare the percentage of diagonal outcomes to the percentages of selfish and cooperative outcomes. We classify the outcome according to which of these three categories has the largest percentage. If the diagonal category has the largest percentage and if neither off-diagonal box occurs more than twice as often as the other, we classify the outcome as alternating. If one diagonal outcome is more than twice as likely, we classify the outcome as asymmetric.

In Fig. 1 we plot those outcomes using barycentric coordinates. We classify (C, S) and (S, C) as *off-diagonal* outcomes which allows us to display the outcomes from each game in a two-dimensional simplex, where the three possible outcomes are *cooperate* (C, C) , *selfish* (S, S) , and *diagonal* (C, S) and (S, C) .

In each game, a joint payoff maximizing outcome is achieved nearly every time. Only in the Knife Edge and Alternation games do we see variation in the outcomes. We can further unpack the outcomes in these two games by computing how often the outcome results in an alternating strategy between the off-diagonals in which each player receives the same average payoff and how often the players settle into an asymmetric equilibrium. Those results are shown in Fig. 2.¹³ In Knife Edge more than half the time the agents are able to learn to alternate. In Alternation they learn to alternate approximately two-thirds of the time.

3.3. Set dependence

Given these baseline results, we next explore the extent to which we see *set dependence*. We explore set dependence in both outcomes and behaviors. We focus here first on the effect of various initial games on two games: *Knife Edge* and *Stag Hunt*. Doing so enables us to demonstrate how and why set dependence of outcomes occurs as well as revealing how behaviors drive that dependence.

We begin by showing the dramatic differences in outcomes that occur in Knife Edge following different initial games in the sequence. Fig. 3 shows the outcomes for the Knife Edge game given different initial games. When Knife Edge is the first game in the sequence, agents playing Knife Edge produce approximately twice as much alternating outcomes as cooperative outcomes, and about three times as many cooperative outcomes as asymmetric outcomes. Thus, alternating outcomes are fifteen times as likely as asymmetric outcomes. In contrast, when the agents play Knife Edge after Bottom Left, asymmetric outcomes become more likely. In fact, the agents arrive at asymmetric outcomes more than six times as often as alternating outcomes. When the agents play Knife Edge after Stag Hunt, Prisoner's Dilemma, or Self Interest, asymmetric outcomes almost never occur. Finally, selfish outcomes prove to be rare occurrences in every setting except for when Knife Edge follows Alternation.

Many of these findings have intuitive explanations. In the Stag Hunt and Prisoner's Dilemma games, agents cooperate and that cooperation bleeds over into the Knife Edge game. Similarly, in Bottom Left, the agents settle on an asymmetric outcome and that outcome often is maintained in Knife Edge. But other results are less intuitive. Why, for example, does playing the Self Interest game first make the agents more likely to alternate than if they had not played any game at all? And why does playing Bottom Left first make the agents more likely to learn to cooperate?

To explain these less intuitive findings, it is useful to consider the behaviors that produce them. When Knife Edge is the initial game, the agents begin with arbitrary strategies and then learn how to play each game. A natural way to represent those strategies relies on how the agents' behavior changes as a function of the agent's own behavior and that of its opponent. Recall that an agent can be thought of as having two mental states: one in which it is being cooperative, and one in which it is being selfish. These mental states correspond to the agents' action in the previous round. Given that action, there exist four possible transitions. First, the agent could remain *fixed*, i.e. keep taking the same action. Second, the agent could *match* the action of its opponent. So for example, an agent using Grim Trigger matches if taking the cooperative action and switches if taking the selfish action whereas an agent playing Tit for Tat matches in each state. Third, the agent could *switch* to the other action. This transition rule would be useful in the Alternation Game. Finally, the agent takes the *opposite* action of its opponent. This last transition rule will be useful in the Top Left and Bottom Right games given that the agents would like to take opposing actions.

In Fig. 4, we plot a randomly drawn initial distribution of strategies using this classification of strategies. The size of the circles corresponds to the propensity of that type of strategy.¹⁴ Thus, the size of the circle that is fixed in the selfish state and matches in the cooperative state corresponds to agents playing Grim Trigger and to agents playing Always Selfish, where the two are differentiated only by their initial states: Grim Trigger automata start out cooperating and Always Selfish start out defecting. Given that we generate these automata randomly, agents' initial strategies are equally likely across the possible classifications. As the agents evolve strategies this distribution will change.

In Fig. 5, we plot the average automata after the agents have learned to play Knife Edge. Three types of strategies tend to evolve: $(match, fixed)$, $(match, match)$, and $(switch, match)$. The first of these pairs encodes *Grim Trigger* (assuming the agents cooperate in the first round, which they do). The second pair can encode either *Tit for Tat* or *Alternating* depending on whether the two agents begin taking the same action or different actions. The third encodes *Alternating*.

In Fig. 6, on the left we plot the distribution over automata after the agents have learned to play Bottom Left, a game that advantages the row player. On the right we plot the distribution for the agents after they have learned to play Knife Edge given that their

¹³ Strategies are classified by modal outcomes. A strategy was classified as asymmetric if one of the off-diagonal outcomes was more than twice as probable as the other.

¹⁴ The distribution of strategies graphs are based on 100 trials. These graphs are used to illustrate strategies.

initial behaviors were those learned for Bottom Left. Notice that when playing Bottom Left, the agents tend to learn a behavior classified as (*opposite, fixed*). This strategy will cooperate when the other agent is being selfish and remain selfish when being selfish. When the Knife Edge game is played first, this pair of strategies rarely occurs. However, often when Knife Edge follows Bottom Left, that behavior remains in place when Knife Edge is played. The fact that when playing Bottom Left first the agents learn to be fixed when in the selfish state also explains why so much cooperation occurs when agents start playing Knife Edge. Grim Trigger (represented as (*match, fixed*)) can induce cooperation by severely punishing those who don't.

A similarly strong behavioral spillover can be seen by comparing what behaviors the agents learn to play in the Prisoners' Dilemma to what they learn to play in Knife Edge after first playing the Prisoners' Dilemma. As shown in Fig. 7, in the Prisoners' Dilemma the agents learn *match* in the cooperative state. In the selfish state, they may play *fixed* (Grim Trigger), *match* (Tit for Tat), *switch* (Punish Once), or *opposite* (Win Stay, Lose Shift). Each of these behaviors produces cooperation. When the agents learn to play Knife Edge after having played the Prisoner's Dilemma, many of these behaviors remain.

A close look at Fig. 7 also reveals why playing Knife Edge following the Prisoners' Dilemma produces almost no asymmetric outcomes. Asymmetric outcomes require agents to be fixed or play the opposite in each state. After playing the Prisoners' Dilemma, almost all of the agents learn to *match* in the cooperative state. Given that mutual cooperation produces a good payoff, and that cooperating when the other is selfish does not, this behavior will only be abandoned if the other player introduces an alternating strategy.

We next demonstrate set dependent outcomes in Stag Hunt. Recall that when Stag Hunt is the first game in the sequence that the agents always learned to cooperate. In Fig. 8, we show the percentage of selfish outcomes for Stag Hunt following each of five games.

When Stag Hunt follows Self Interest, approximately ten percent of the time the agents remain selfish. This is to be expected. In the Self Interest game the agents learn to be selfish and this behavior is not always unlearned in Stag Hunt. But, quite surprisingly, when Stag Hunt follows Bottom Left of Top Right, thirty percent of the time the agents learn to be selfish in Stag Hunt, a nearly three-fold increase over the level of selfish outcomes following Self Interest. To see why, we need only look at the behaviors in the games.

Fig. 9 shows behavior in Stag Hunt when it is the first game in the sequence. Agents learn to play either be *fixed* or *match* in the cooperative state. It doesn't really matter what they do in the selfish state as that is not reached.

Fig. 10 shows behavior following Bottom Left and behavior following Self Interest. Playing Bottom Left, the agents often learn *opposite* in the cooperative state. In effect, they learn to take turns, to do the opposite of what the other player does. So, if the other player cooperates the agent will then be selfish. This behavior produces selfish behavior in the Stag Hunt game. In contrast, when agents learn the Self Interest game, they learned *fixed* behavior in the selfish state but they're no more likely to learn *opposite* behavior in the cooperative state than any other behavior. This makes agents *more likely* to learn to be cooperative following self interest rather than less likely.

A similar intuition explains why there exists more selfish play in Stag Hunt following Alternation. Agents who first play Alternation often learn either *opposite* or *switch* behavior in the cooperative state. (See Fig. 11). Switching to selfish behavior immediately after cooperating makes it more difficult for cooperation to emerge.

3.4. Path dependence

The previous results demonstrated set dependence of outcomes and of behaviors. We now explore the extent to which our model produces path dependence. Path dependence exists if the same games played in a different order produced distinct outcomes and behaviors. We begin our analysis by considering two pairs of games that include the Alternation game and another game. We first analyze Knife Edge and Alternation. The second pair consists of Stag Hunt and Alternation.

In Fig. 12, we show outcomes in Knife Edge and in Alternation when Knife Edge is played first and when Alternation is played first. The graph on the top shows outcomes for Knife Edge. When agents play Knife Edge second, they are far less likely to cooperate and are more likely to alternate or to act selfishly. The graph on the bottom shows that when Alternation is played second, agents are much more likely to alternate.

These may seem like minor differences, but notice that when Alternation is played first, more selfish behavior occurs in both games, resulting in lower payoffs than if Alternation is played second. Therefore, it is better to play Knife Edge first and then Alternation rather than playing Alternation first and then Knife Edge. The reason for this is that Alternation produces some selfish behavior which then gets learned in Knife Edge. When Knife Edge is played first, the agents never learn to be selfish which reduces the amount of selfish behavior in Alternation. This can be seen at an even deeper level by referring back to Fig. 5. Notice that when Knife Edge is played first, the agents almost never learn the opposite behavior.

Next, we consider the pair of games Alternation and Stag Hunt. Fig. 13 shows the outcomes in Stag Hunt and Alternation game when Stag Hunt is placed first in the sequence and when it is placed second. When Stag Hunt is played first, the agents are much more likely to learn to be cooperative than if Stag Hunt is introduced after the Alternation game. When Alternation is played first, the agents sometimes learn to *switch* when in the cooperative state. This behavior is difficult to unlearn and results in selfish behavior in Stag Hunt. This unfortunate spillover is an argument for playing Stag Hunt first. If Stag Hunt is played first, the agents are also more likely to alternate and less likely to be selfish in the Alternation Game. So, once again, we see path dependence in the performance of institutions.

4. Institutional path dependence

Having demonstrated path dependence of both behaviors and outcomes, we now extend the analysis and discuss how behavioral spillovers might produce institutional path dependence. Rather than choosing actions, our analysis expands to consider choice over institutions, proposed as ways to frame interaction to overcome new problems. In the previous section, we have shown that the performance of an institution can depend on the previous institutions because of behavioral spillovers. If a particular institution doesn't

Table 1
Percent time the equilibrium outcome is located when each game is played alone.

Game	Equilibrium outcome	% efficient
BL	Off-diagonal	0.968
TR	Off-diagonal	0.963
SI	Selfish	1.00
SH	Cooperate	0.985
PD	Cooperate	0.998
KE	Cooperate or alternate	0.353 (C) or 0.560 (ALT)
AL	Alternate	0.680

perform well following a prior institution or set of institutions then we can assume that the institution won't be chosen. This could occur because the society has the foresight to know it won't function well or because the institution is tried, fails, and is replaced with one that performs better.

The existence of two previous games implies that agents have a larger repertoire of behaviors from which to choose an initial strategy. Here, we will assume that in the third game, each agent randomly chooses either its strategy in the first game or its strategy in the second game.¹⁵ This construction assumes that the agents see all pairs of games as equally similar—admittedly a strong assumption, but a reasonable starting point.

We explore the extent to which institutions are *path limiting*, defined as the percentage of paths of institutional choices that are inefficient. To make this assumption formal, we will assume that if given an institutional context the efficient equilibrium is located five percent less often than if the game were played alone, then the game will not be chosen. Formally, we will say that a path is *inefficient* if efficiency falls by at least five percent. For example, in the Knife Edge game the efficient equilibrium is almost always learned. However if Knife Edge follows the Alternation game, then the selfish outcome is learned more than five percent of the time. We therefore will assume that if the Alternation game is chosen first that Knife Edge will not be chosen second.

With seven possible games there exist 343 possible sequences of length three. Given an initial game, there are 49 possible two-game continuation sequences. In Fig. 14, we show all possible sequences following the Prisoners' Dilemma game.

Not all of these paths may be efficient. In Fig. 15, we erase the inefficient paths following the Bottom Left Game based on 400 trials of every three-game sequence. What's clear from the figure is that Bottom Left *limits* the possible institutional sequences.¹⁶

In Fig. 16, we show the efficient paths following the Prisoners' Dilemma game.¹⁷ Notice that many more paths are possible. The Bottom Left game is much more path-limiting than the Prisoners' Dilemma. The reason for this is that the Prisoners' Dilemma produces diverse behaviors in the selfish state. That diversity enables the agents to learn the efficient equilibrium for other games. They're less likely to be stuck in the basin of a bad equilibrium. In contrast, in Bottom Left, agents learn similar behaviors. This increased behavioral coherence in each state hinders exploration. Thus, we find value in diverse behaviors within a game because it aids learning.

The assumption that agents initially choose the behavior of only the nearest game is a strong assumption. People might well differ in which games they believe to be near to one another leading them to choose different initial behaviors. This diversity of behaviors from which to choose will also increase the likelihood of more efficient equilibria (at least in the games that we consider here). Therefore, ideal sequences of institutions should produce diverse behavior both within and across games. This second intuition echoes the Bednar and Page (2014) result that early institutions should be diverse to build up diverse repertoires.

5. Discussion

Within our model we demonstrate an interplay between the set and sequence of institutions, behavior, institutional outcomes, and optimal institutional choice. Though our model relies on a particular form of behavioral spillovers to produce institutional path dependence, any of a variety of spillovers should produce similar results. Behavioral spillovers are a possible additional cause of institutional path dependence — one that complements existing theories based on increasing returns and negative externalities (Pierson, 2000, 2004; Page, 2006). For example, what Pierson describes as increasing returns often includes institutions that leverage similar behavior. And, the path limiting behavior that we identify in our model can be interpreted as a type of negative externality.

Our findings also align with the idea that past experience affects the performance of new institutions, and, by extension, should be taken into account when designing or choosing among institutions. In demonstrating that institutions reinforce one another by relying on similar behavioral repertoires, we link institutional choices with some dimensions of culture. Institutions have considerable influence over behaviors, norms, and culture, so such a link makes intuitive sense.

The empirical question of how far back one must look to identify the relevant spillovers is an interesting question. In comparing the experiences of transitioning to democracy in Kazakhstan, Kyrgyzstan, and Uzbekistan, Jones Luong (2002) finds that they did not revert to historical clan-based representational systems. All chose regionally-based electoral systems that mirrored the Soviet system put in place by Gorbachev as a result of his policies of Glasnost and Perestroika. Thus, to quote Putnam (1993), we may need not reach all the way back to the *mists of the dark ages* to explain behavior and institutional choice.

¹⁵ An alternative assumption would be that the agents choose the behavior from the game that's most similar to the new game. See Bednar and Page (2014).

¹⁶ Please see Tables 1, 3 and 5 for the statistics behind Figure 15.

¹⁷ Please see Tables 1, 2 and 4 for the statistics behind Figure 16.

Table 2
Two game sequences: Bottom Left.

Game 1	Game 2	% efficient outcome game 2
BL	TR	0.848 ^a
BL	SI	1.00
BL	SH	0.738 ^a
BL	PD	0.905 ^a
BL	KE	0.360 (C); 0.495 (ALT)
BL	AL	0.055 ^a

^a Efficiency decreases by more than 5% compared to game played alone.

^b Efficiency increases by more than 5% compared to game played alone.

Table 3
Two game sequences: Prisoners' Dilemma.

Game 1	Game 2	% efficient outcome Game 2
PD	BL	0.745 ^a
PD	TR	0.72 ^a
PD	SI	1.00
PD	SH	1.00
PD	KE	0.553 (C) ^b ; 0.415 (ALT)
PD	AL	0.823

^a Efficiency decreases by more than 5% compared to game played alone.

^b Efficiency increases by more than 5% compared to game played alone.

The unique contribution of this paper is to focus on microlevel behavior and how it transfers across institutions. This approach opens up the possibility of a building a methodological bridge between those searching for culture-free regularities and those who favor exceptionalism and treat each situation as unique. The tension between universalistic rational choice models and area studies animates multiple conversations within political science and economics. One route to compromise is to divide up the various dimensions of analysis. Some outcomes—say, the number of political parties—might be thought best explained by rational choice models and other outcomes, such as rates of corruption, might be thought to require an appeal to culture. Models like ours do not require dividing up the domains of analysis. They provide a potential explanation that is consistent with both camps. Behavior may well be purposive but also influenced by the cultural context. Outcomes that are consistent across cultures may be equilibria that have greater immunity to behavioral spillovers.

We in no way mean to suggest behavior as the only candidate explanation. Models based on beliefs or belief systems that drive disparate institutional performances can also produce forms of historical dependence.¹⁸ To unpack the differences between our *behavior-centric* approach and *belief-centric* models requires explicit characterization of behaviors and beliefs. Multiple techniques exist to elicit beliefs. What is less widely known is that there exists a large body of work that identifies behavioral rules, such as studies that show how citizens use rules of thumb when they make inferences about new political candidates (Mondak and McCurley, 1994; Ottati, 1990; Brent and Granberg, 1982)¹⁹ or how jurors use heuristics to weigh evidence in a trial (Saks and Kidd, 1980–1981).

So while numerous studies show that people rely on cognitive “shortcuts” both in decision contexts and in strategic setting (see Kahneman and Tversky, 1979; Camerer, 2003) and how such heuristics can be time- and energy-saving and produce nearly optimal actions (Gigerenzer and Selten, 2002; Clark, 1997; Barber, 1977), less has been done to analyze whether those behaviors spill into other contexts. Exceptions include experimental work (Cason et al., 2012) and some case studies. In addition to the experimental work of Bednar et al. (2012a,b), Simpson (2004) finds that cooperation is ubiquitous among citizens who apply prosocial heuristics to social dilemma games such as the Prisoner's Dilemma, but is not common among citizens who are individualistic.

Case study support for behavioral spillovers includes research on interethnic conflict and cooperation by Arfi (2000) who demonstrates that interethnic cooperation emerges and persists even when agents care little about the future and lack punishment mechanisms. Relatedly, Fearon and Laitin (1996) show that interethnic cooperation, once established, is likely to persist even after institutional mechanisms for cooperation are removed. These results demonstrate the potential explanatory power of moving beyond isolated, single game analyses of behavior and outcomes without necessarily supporting any one type of behavioral spillover over another.

To the extent that one views behavioral regularities as a part of culture, our results engage ethnographic and case study research on the influence of culture on institutional and organizational outcomes (Platteau, 2000). On a range of policy problems, such as environmental protection, state formation, and economic development, it is well documented that one-size-fits-all solutions often fail (Ostrom et al., 2007) and sometimes do so miserably (Stiglitz, 2003, Easterly 2001, Dollar and Svensson, 2000). Further, wholesale

¹⁸ See Greif (1994, 2006), Greif and Laitin (2004), and Putnam (1993).

¹⁹ Brady and Sniderman (1985) find that the most common heuristic used by Americans is what they call the “likability heuristic” — that is, Americans give much more weight to how much they personally like each candidate than to other characteristics like policy positions or background when making choices about whom to support.

Table 4
Three game sequences: Bottom Left.

Game 1	Game 2	Game 3	% efficient outcome Game 3
BL	TR	SI	1.00
BL	TR	SH	0.640 ^a
BL	TR	PD	0.810 ^a
BL	TR	KE	0.203 ^a (C); 0.233 ^a (ALT)
BL	TR	AL	0.198 ^a
BL	SI	TR	0.938
BL	SI	SH	0.820 ^a
BL	SI	PD	0.933 ^a
BL	SI	KE	0.333 (C); 0.223 (ALT)
BL	SI	AL	0.215 ^a
BL	SH	TR	0.888 ^a
BL	SH	SI	1.00
BL	SH	PD	0.955
BL	SH	KE	0.335 (C); 0.250 (ALT)
BL	SH	AL	0.073 ^a
BL	PD	TR	0.910 ^a
BL	PD	SI	1.00
BL	PD	SH	0.953
BL	PD	KE	0.460 ^b (C); 0.380 (ALT)
BL	PD	AL	0.068 ^a
BL	KE	TR	0.713 ^a
BL	KE	SI	1.00
BL	KE	SH	0.888 ^a
BL	KE	PD	0.950
BL	KE	AL	0.230 ^a
BL	AL	TR	0.868 ^a
BL	AL	SI	1.00
BL	AL	SH	0.273 ^a
BL	AL	PD	0.900 ^a
BL	AL	KE	0.313 (C); 0.158 (ALT)

^a Efficiency decreases by more than 5% compared to game played alone.

^b Efficiency increases by more than 5% compared to game played alone.

copying of institutions that have worked elsewhere rarely works in new settings (Fish, 1995). The failure of the Washington Consensus policies to deliver economic growth in developing countries in the 1990s is representative of this larger point. Some countries, such as Botswana in recent years, have responded well to World Bank and IMF interventions and exhibited rapid growth, while others, such as Zambia, have seen few positive gains, and even exhibited negative growth (Irving, 1998).

Evidence that things don't go as planned could be explained with any of a variety of competing models. Rigorous empirical testing of our model lies beyond the scope of the current paper, but we can suggest some cases that warrant deeper investigation. These include the efforts Chinese Government in the early 1950s to implement gender equality in rural households. Hershatter (2002) finds that those that proved most enduring were those that validated existing divisions of labor. Allen (2006), on the other hand, finds that laws that ban arranged marriage and bride-buying in rural China did not translate into better treatment of women. Those laws had few existing behaviors to leverage.

Relatedly, Karl (1997) has suggested that oil-rich states have trouble developing diverse markets because the benefits that accrue from oil hinder the ability of markets to create incentives to innovate. Grzymala-Busse (2010) has found significant variation in the performance of stock markets in Eastern Europe based upon the timing of the introduction of regulatory institutions relative to stock markets. When market behavior was allowed to develop in the absence of regulation, regulation was much less effective once introduced than if behavior had developed with regulatory oversight from the start. Greif (1994, 2006) shows that institutional performance depends on the characteristics of societies — the trust-based, segregated economic relations of the 11th century Maghribi worked well as long as the trading circle was small, but the individualistic 12th century Genovese had institutions in place to enforce contracts, giving them the advantage in long-distance trading.

In all of these cases, explanations for the success and failure of institutions share an emphasis on the role historically contingent behaviors. We do not mean to imply that these prove the veracity of our model, but to suggest the potential value in trying to measure these behavioral spillovers empirically. Do they take the form we assume here—by constraining initial behavior—or do they operate in some other way?

Finally, although our analysis and discussion have focused on societal-level institutional effects, we might have interpreted our entire analysis in the context of organizational culture. That literature provides many examples of path dependent behavior including strong evidence of imprinting in which early actions taken by a leader or focal actor persist even though the conditions under which they emerge no longer hold. Imprinting has been found in craft unions, department stores, banks, newspapers, and high-tech firms (Stinchcombe, 1965; Swaminathan, 1996; Boeker, 1989). Within organizations, imprinting applies to routines and learning rules (Cohen and Bacdayan, 1994).

Table 5
Three game sequences: Prisoners' Dilemma.

Game 1	Game 2	Game 3	% efficient outcome Game 3
PD	BL	TR	0.720 ^a
PD	BL	SI	1.00
PD	BL	SH	0.993
PD	BL	KE	0.453 (C) ^b ; 0.440 (ALT)
PD	BL	AL	0.325 ^a
PD	TR	BL	0.678 ^a
PD	TR	SI	1.00
PD	TR	SH	0.993
PD	TR	KE	0.468 ^a (C) ^b ; 0.418 (ALT)
PD	TR	AL	0.303 ^a
PD	SI	BL	0.930
PD	SI	TR	0.915
PD	SI	SH	1.00
PD	SI	KE	0.463 (C) ^b ; 0.500 (ALT)
PD	SI	AL	0.820 ^b
PD	SH	BL	0.763 ^a
PD	SH	TR	0.165 ^a
PD	SH	SI	1.00
PD	SH	KE	0.485 (C) ^b ; 0.450 (ALT)
PD	SH	AL	0.868 ^b
PD	KE	BL	0.345 ^a
PD	KE	TR	0.710 ^a
PD	KE	SI	1.00
PD	KE	SH	0.998
PD	KE	AL	0.930 (C) ^b ; 0.070 (ALT)
PD	AL	BL	0.785 ^a
PD	AL	TR	0.835 ^a
PD	AL	SI	1.00
PD	AL	SH	1.00
PD	AL	KE	0.243 (C); 0.713 (ALT) ^b

^a Efficiency decreases by more than 5% compared to game played alone.

^b Efficiency increases by more than 5% compared to game played alone.

Societal level imprinting can also occur. Ebay's method of auctioning off goods has become prevalent in other markets because so many people have evolved strategies for those rules. Revenue equivalence of auction designs notwithstanding, it could well be that in other environments other auction designs might work better than Ebay's rules.²⁰

As Clark (1997) reminds us, "nature is heavily bound by achieved solutions to previously encountered problems" (p. 81). Thus, when choosing or designing an institution or an institution, we should not limit attention to the equilibria it implements. We ought also consider existing individual and collective behavioral repertoires, and take into account their effects on institutional performance, choice, and even design.

Acknowledgments

We thank Ron Harstad, Anna Grzymala-Busse, Bill Zimmerman and Gerry Mackie for the helpful conversations, and seminar participants at ICER (Turin) and APSA, especially our discussants Daniel Diermeier and Tasos Kalandrakis. The paper benefited from feedback from participants at the Venice Summer Institute Workshop on Behavioral Political Economy. For research assistance, we are indebted to Jennifer Miller and Neill Mohammad. Funds to support this research project were provided by the McDonnell Foundation, an NSF IGERT grant, AFOSR-MURI Grant 57100001867, and the U.S. Army Research Laboratory of the U.S. Army Research Office under contract/grant number W911NF1010379.

Appendix A. Mathematical proofs

Proof of Claim 1. For Top Right, Bottom Left, and Self Interest, let each automaton begin by taking the action that produces the highest joint payoff and continues to take that action. Any deviation produces a lower payoff. This completes the proof for those games. For Knife Edge, Stag Hunt and the Prisoners' Dilemma, assume that every agent plays *Tit for Tat* and begins by cooperating. Given the length of memory of the automata, any automata that does not always cooperate (and therefore receive the same payoff) must deviate in either the first, second, or third round. That deviation will be followed by a deviation from *Tit for Tat* so the average

²⁰ Holbrook et al. (2000) similarly notes that visions for the future can be limited by past experiences. A vision differs from a routine or learning rule, but the effects could be related.

payoff will be lower for the automata not playing *Tit for Tat*. For Alternation, assume that every agent uses an automata that matches in each state and that row playing agents initially cooperate and column player agents initially play choose selfish. Any deviation that does not produce the same sequence of outcomes must produce an outcome of either (C, C) or (S, S) in one round. Without loss of generality, assume the deviant strategy is the row player. Given that the other agents all match, any (C, C) must follow a (C, S) and any (S, S) must follow an (S, C). The first possibility creates the sequence of outcomes (C, S), (C, C), (S, C), (C, S), (C, C), (S, C) which results in a lower average payoff for the first player. The second produces the outcome sequence (S, C), (S, S), (C, S), (S, C), (S, S), (C, S) which also has a lower average payoff, completing the proof.

Proof of Claim 2. Recall that each agent plays its two neighbors to the left and right. We will assume that there exist ten agents and that agents one through five play an efficient strategy one and that agents six through ten play an inefficient strategy. We will refer to these as the *efficient agents* and the *inefficient agents*. This construction extends to the case of more agents trivially as more agents merely implies more agents in the interior of the efficient and inefficient regions. Note that by symmetry, we only need consider the payoffs for agents three through eight. Agents three through five play the efficient strategy and agents six through eight play the inefficient strategy. It suffices to show that that no efficient agent would like to switch to the inefficient strategy and that inefficient agent would like to switch to the efficient strategy. The following two conditions are sufficient for a Best Neighbor Equilibrium.

Efficient stability. *Agent three earns a higher payoff than either agent six or agent seven.*

Inefficient stability. *The maximum payoff from agents six through eight exceeds the maximum payoff of agents four and five.*

We prove the result by constructing strategies for each game.

A.1. Self Interest

Let the efficient agents play *All Selfish* and the inefficient agents play *Silly Grim*: {C, C, S, C, S, S}. *Silly Grim* cooperates in the first round and continues to cooperate as long as the other agent is selfish. If the other agent cooperates, then the agent switches to selfish and plays selfish forever. Two *Silly Grim* agents each earn eight in every round but the first. When a *Silly Grim* meets an *All Selfish*, the *Silly Grim* earns zero and the *All Selfish* earns six. Let M equal the number of rounds. Payoffs for agents three through eight are given in the table below. We assume each agent plays each other agent twice. Once as the row player and once as the column player.

Agent 3 <i>All Selfish</i>	Agent 4 <i>All Selfish</i>	Agent 5 <i>All Selfish</i>	Agent 6 <i>Silly Grim</i>	Agent 7 <i>Silly Grim</i>	Agent 8 <i>Silly Grim</i>
$64M$	$60M$	$56M$	$32(M - 1)$	$48(M - 1)$	$64(M - 1)$

By inspection, payoffs satisfy both efficient and inefficient stability.

A.2. Top Right (Bottom Left)

Let the efficient agents always play outcomes in the Top Right (TR). Define the strategy *Bottom Left Once (BL1)* as follows: a row agent initially plays S and a column agent initially plays C. The row agent matches its opponent when playing S and remains fixed when playing C. The column agent matches its opponent when playing C and remains fixed once it plays S. *Note that even though playing S is a dominant strategy for the column agent, if playing against a row agent using Bottom Left Once, the column agent receives a higher long term payoff by playing C in the first round to lock the row agent into playing C.* Otherwise, the row agent continues to play S.

Agent 3 TR	Agent 4 TR	Agent 5 TR	Agent 6 BL1	Agent 7 BL1	Agent 8 BL1
$64M$	$60M$	$56M$	$32(M - 1)$	$48(M - 1)$	$64(M - 1)$

By inspection, payoffs satisfy efficient and inefficient stability. A symmetric construction proves the result for Bottom Left.

A.3. Prisoners' Dilemma and Stag Hunt

Let the efficient agents use *Tit for Tat (TFT)* and the inefficient use the strategy *Selfish Switch Match (SSM)* which is initially selfish. After being selfish it cooperates and after cooperating it matches the behavior of its opponent. Two *Selfish Switch Match* players will cooperate in every round but the first and a *Selfish Switch Match* and a *Tit for Tat* will alternate on the off diagonals. We show the payoffs for the Prisoners' Dilemma. The payoffs for Stag Hunt also satisfy the two sufficient condition.

Agent 3 TFT	Agent 4 TFT	Agent 5 TFT	Agent 6 SSM	Agent 7 SSM	Agent 8 SSM
64M	59M	54M	54M – 20	59M – 30	64M – 40

By inspection, payoffs satisfy efficient and inefficient stability.

A.4. Alternation

Let the efficient agents use *Alternate* (ALT), which initially cooperates if the row player and is selfish if the column player and then switches in each state. Let the inefficient players play *Always Selfish* (AS). We first show the payoffs for the Alternation game.

Agent 3 ALT	Agent 4 ALT	Agent 5 ALT	Agent 6 AS	Agent 7 AS	Agent 8 AS
32M	24M	16M	32M	24M	16M

By inspection, the payoffs satisfy efficient and inefficient stability.

Knife Edge lets the efficient agents use *Grim Trigger* and the inefficient players play *Selfish Switch Match*.

Agent 3 GT	Agent 4 GT	Agent 5 GT	Agent 6 SSM	Agent 7 SSM	Agent 8 SSM
64M	56M + 16	48M + 32	48M + 24	56M + 8	64M – 8

By inspection, the payoffs satisfy efficient and inefficient stability.

References

- Akerlof, G., Shiller, R., 2009. *Animal Spirits: How Human Psychology Drives the Economy, and Why It Matters for Global Capitalism*. Princeton University Press, Princeton, NJ.
- Arfi, B., 2000. "Spontaneous" Interethnic Order: The Emergence of Collective, Path-Dependent Cooperation. *Int. Stud. Q* 44 (4), 563–590.
- Barber, G.M., 1977. Sequencing highway network improvements: a case study of South Sulawesi. *Econ. Geogr.* 53, 55–69.
- Bednar, J., Page, S.E., 2007. Can game(s) theory explain culture? The emergence of cultural behavior within multiple games. *Ration. Soc.* 19, 65–97.
- Bednar, J., Page S.E., 2014. When order affects performance: Institutional sequencing, cultural sway, and behavioral path dependence. Unpublished manuscript.
- Bednar, J., Bramson, A., Jones-Rooy, A., Page, S.E., 2010. Emergent cultural signatures and persistent diversity: a model of conformity and consistency. *Ration. Soc.* 22, 407–444.
- Bednar, J., Page, S.E., Toole, J., 2012a. Revised-path dependence. *Polit. Anal.* 20, 146–156.
- Bednar, J., Chen, Y., Liu, T.X., Page, S.E., 2012b. Behavioral spillovers and cognitive load in multiple games: an experimental study. *Games Econ. Behav.* 74, 12–31.
- Bendor, J., Diermeier, D., Ting, M., 2003. A behavioral model of turnout. *Am. Pol. Sci. Rev.* 97, 261–280.
- Boeker, W., 1989. Strategic change: the effects of founding and history. *Acad. Manag. J.* 32, 484–515.
- Brady, H.E., Sniderman, P.M., 1985. Attitude attribution: a group basis for political reasoning. *Am. Polit. Sci. Rev.* 79, 1061–1078.
- Brent, E., Granberg, D., 1982. Subjective agreement with the presidential candidates of 1976 and 1980. *J. Pers. Soc. Psychol.* 42, 393–403.
- Camerer, C., 2003. *Behavioral Game Theory: Experiments on Strategic Interaction*. Princeton University Press, Princeton, NJ.
- Cason, T.N., Savikhin, A.C., Sheremeta, R.M., 2012. Behavioral spillovers in coordination games. *Eur. Econ. Rev.* 56, 233–245.
- Clark, A., 1997. *Being There: Putting Brain, Body, and World Together Again*. MIT Press, Cambridge, MA.
- Cohen, M.D., Bacdayan, P., 1994. Organizational routines are stored as procedural memory: evidence from a laboratory study. *Organ. Sci.* 5, 554–568.
- Diermeier, D., Krehbiel, K., 2003. Institutionalism as methodology. *J. Theor. Polit.* 15, 123–144.
- Dolan, P., Galizzi, M.M., 2015. Like ripples on a pond: behavioral spillovers and their consequences for research and policy. *J. Econ. Psychol.* 47, 1–16.
- Dollar, D., Svensson, J., 2000. What explains the success or failure of structural adjustment programmes? *Econ. J.* 110, 894–917.
- Fearon, J., Laitin, D., 1996. Explaining interethnic cooperation. *Am. Polit. Sci. Rev.* 90, 715–735.
- Fernández, R., Fogli, A., 2006. Fertility: the role of culture and family experience. *J. Eur. Econ. Assoc.* 4, 552–561.
- Fish, M.S., 1995. *Democracy from Scratch: Opposition and Regime in the New Russian Revolution*. Princeton University Press, Princeton, NJ.
- Fudenberg, D., Levine, D., 1998. *The Theory of Learning in Games*. MIT Press, Cambridge, MA.
- Gigerenzer, G., Selten, R. (Eds.), 2002. *Bounded Rationality: The Adaptive Toolbox*. MIT Press, Cambridge, MA.
- Gilboa, I., Schmeidler, D., 1995. Case-based decision theory. *Q. J. Econ.* 110, 605–639.
- Golman, R., Page, S.E., 2010. Basins of attraction and equilibrium selection under different learning rules. *Evol. Econ.* 20, 49–72.
- Greif, A., 1994. Cultural beliefs and the organization of society: a historical and theoretical reflection on collectivist and individualist societies. *J. Polit. Econ.* 102, 912–950.
- Greif, A., 2006. *Institutions and the Path to the Modern Economy*. Cambridge University Press, Cambridge, UK.
- Greif, A., Laitin, D., 2004. A theory of endogenous institutional change. *Am. Polit. Sci. Rev.* 98, 633–652.
- Grzymala-Busse, A., 2010. The best laid plans: the impact of informal rules on formal institutions in transitional regimes. *Stud. Comp. Int. Dev.* 45, 311–333.
- Henrich, J., Boyd, R., Bowles, S., Gintis, H., Fehr, E., Camerer, C. (Eds.), 2004. *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence in Fifteen Small-Scale Societies*. Oxford University Press, Oxford, UK.
- Henrich, J., Ensminger, J., McElreath, R., Barr, A., Barrett, C., Bolyanatz, A., Cardenas, J.C., Gurven, M., Gwako, E., Henrich, N., Lesorogol, C., Marlowe, F., Tracer, D., Ziker, J., 2010. Markets, religion, community size, and the evolution of fairness and punishment. *Science* 327, 1480–1484.
- Hershatter, G., 2002. The gender of memory: rural Chinese women and the 1950s. *Signs J. Women Cult. Soc.* 28, 43–70.

- Holbrook, D., Cohen, W.M., Hounshell, D.A., Klepper, S., 2000. The nature, sources, and consequences of firm differences in the early history of the semiconductor industry. *Strateg. Manag. J.* 21, 1017–1041.
- Hurwicz, L., 1972. On informationally decentralized systems. In: McGuire, C.B., Radner, R. (Eds.), *Decision and Organization: A Volume in Honor of Jacob Marschak*. North-Holland Publishing Company, Amsterdam, pp. 297–336.
- Hurwicz, L., Schmeidler, D., 1978. Construction of outcome functions guaranteeing existence and Pareto-optimality of Nash equilibria. *Econometrica* 46, 1447–1474.
- Kahneman, D., Tversky, A., 1979. Prospect theory: an analysis of decision under risk. *Econometrica* 47, 263–292.
- Kalai, E., Stanford, W., 1988. Finite rationality and interpersonal complexity in repeated games. *Econometrica* 56, 397–410.
- Karl, T., 1997. *The Paradox of Plenty: Oil Booms and Petro-States*. University of California Press, Berkeley, CA.
- Medin, D., Bennis, W., Chandler, M., 2010. The home-field disadvantage. *Perspect. Psychol. Sci.* 5, 708–713.
- Miller, J.H., Page, S.E., 2007. *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton University Press, Princeton, NJ.
- Mondak, J.J., McCurley, C., 1994. Cognitive efficiency and the congressional vote: the psychology of coattail voting. *Pol. Res. Q.* 47, 151–175.
- Morton, R., Williams, K., 2010. *Experimental Political Science and the Study of Causality: From Nature to the Lab*. Cambridge University Press, Cambridge, UK.
- Nisbett, R.E., Peng, K., Choi, I., Norenzayan, A., 2001. Culture and systems of thought: holistic versus analytic cognition? *Psychol. Rev.* 108, 291–310.
- Ostrom, E., 2005. *Understanding Institutional Diversity*. Princeton University Press, Princeton, NJ.
- Ostrom, E., Janssen, M.A., Anderies, J.M., 2007. Going beyond panaceas. *Proc. Natl. Acad. Sci.* 104, 15176–15178.
- Ottati, Victor C., 1990. Determinants of political judgments: the joint influence of normative and heuristic rules of inference. *Polit. Behav.* 12, 159–179.
- Page, S.E., 2006. Essay: path dependence. *Q. J. Pol. Sci.* 1, 87–115.
- Page, S.E., 2012. A complexity perspective on institutional design. *Polit. Philos. Econ.* 11, 5–25.
- Pierson, P., 2000. Path dependence, increasing returns, and the study of politics. *Am. Pol. Sci. Rev.* 94, 251–267.
- Pierson, Paul, 2004. *Politics in Time: History, Institutions, and Social Analysis*. Princeton University Press, Princeton, NJ.
- Platteau, J.P., 2000. *Institutions, Social Norms and Economic Development*. Routledge Press, New York, NY.
- Putnam, R.D., 1993. *Making Democracy Work: Civic Traditions in Modern Italy*. Princeton University Press, Princeton, NJ.
- Reiter, S., 1977. Information and performance in the (new)² welfare economics. *Am. Econ. Rev.* 67, 226–234.
- Roth, A.E., Prasnikar, V., Okuno-Fujiwara, M., Zamir, S., 1991. Bargaining and market behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: an experimental study. *Am. Econ. Rev.* 81, 1068–1095.
- Rubinstein, A., 1986. Finite automata play the repeated prisoner's dilemma. *J. Econ. Theory* 39, 83–96.
- Saks, M.J., Kidd, R.F., 1980–1981. Human information processing and adjudication: trial by heuristics. *Law Soc. Rev.* 15, 123–160.
- Simpson, B., 2004. Social values, subjective transformations, and cooperation in social dilemmas. *Soc. Psychol. Q.* 67, 385–395.
- Smith, V., 2003. Constructivist and ecological rationality in economics. *Am. Econ. Rev.* 93, 465–508.
- Stiglitz, J.E., 2003. *Globalization and Its Discontents*. W.W. Norton and Company, New York, NY.
- Stinchcombe, A.L., 1965. Social structure and organizations. In: March, J.G. (Ed.), *Handbook of Organizations*. Rand McNally, Chicago, IL, pp. 229–259.
- Swaminathan, A., 1996. Environmental conditions at founding and organizational mortality: a trial by fire model. *Acad. Manag. J.* 39, 1350–1385.
- Tsebelis, G., 1990. *Nested Games: Rational Choice in Comparative Politics*. University of California Press, Berkeley, CA.
- Vriend, N.J., 2000. An illustration of the essential difference between individual and social learning, and its consequences for computational analyses. *J. Econ. Dyn. Control.* 24, 1–19.