

Dynamic coordination with switching costs

Joe S. Ballard* Luke Boosey†

Abstract

We study dynamic coordination with costly switching, where players are initially uncertain regarding *how to coordinate*. We model two players searching for compatible platforms in the presence of two kinds of frictions: noisy outcomes and costly switching. The model predicts that players remain on their current platforms whenever their common belief is sufficiently optimistic; otherwise, they randomize their actions, switching with higher probability whenever they are more pessimistic about compatibility, switching is less costly, or success conditional on compatibility is more likely. In an experiment, we vary (i) whether success is Deterministic or Stochastic when players are on compatible platforms and (ii) the cost of switching. We allow subjects to choose explicit mixed strategies to provide a better test of the theory. The data are consistent with most comparative statics predictions, particularly in the Deterministic case, but differ from some of the other theoretical predictions. To explain the disparities between theory and the data, particularly when learning is noisy and switching is costly, we discuss the impacts of several factors, including cognitive ability, inattention, sensitivity to recent failures, and experience.

Keywords: coordination, switching costs, organizations, experiment, dynamics

JEL: C72, C73, C92, D23, D83

1 Introduction

Coordination is a crucial part of productive efficiency in organizations. Success or productivity may depend on whether or not different units within an organization are able to successfully coordinate their activities, protocols, or procedures. There is by now an enormous literature devoted to the study of coordination games and their applications in a diverse range of social and economic settings. Much of this literature has used experimental methods to explore the power of focal points and the importance of institutional features such as leadership or communication for establishing and maintaining coordination (see, e.g., [?](#), for a review).

In this paper, we study behavior in a dynamic setting where agents are initially uncertain regarding *how to coordinate*. Such a setting may arise in organizations that consist of different units with disparate specializations, norms, or conventions. The source of the initial uncertainty

*Department of Economics, Skidmore College, 815 N. Broadway, Saratoga Springs, NY 12866, USA; E-mail: jballard@skidmore.edu

†Department of Economics, Florida State University, 113 Collegiate Loop, Tallahassee, FL 32306-2180, USA; E-mail: lboosey@fsu.edu.

Acknowledgments. Our sincere thanks go to the editors and two anonymous reviewers for their helpful suggestions and comments during the review process. We gratefully acknowledge funding from the Council on Research & Creativity (CRC) at Florida State University (Planning Grant #044611). This research was conducted with approval from the FSU IRB (STUDY00001674).

may be that the agents lack a common understanding of the strategic environment—in game-theoretic terms, they may lack a common description of the game. Furthermore, the uncertainty may persist if outcomes are noisy when actions are coordinated. Consider, for example, two agents (organizational units) and suppose that each agent knows there are different ways to coordinate but is not well-versed enough in the other agent’s expertise to discern which pairs of activities are compatible. Compatibility in the agents’ actions increases (perhaps stochastically) the output and corresponding benefits of their interaction.

The idea that language barriers or the lack of a common description may impede coordination within organizations has been considered since as far back as ?. For a notable example, ? recount the challenges faced by scientists and engineers working on DNA sequencing at the Broad Institute in Cambridge, Massachusetts. Communication between researchers in different areas was inhibited by the fact that they could not understand each other’s specialized language. In fact, ? also point to this case as an example of an organization that invested specifically in developing a common language to improve coordination and efficiency. Other examples are provided by ?, who attributes the failure of some computer disk drive manufacturing companies to communication issues between the engineering and marketing departments; ?, who notes that engineers and assemblers in a semiconductor manufacturing company lacked a shared understanding of the product and manufacturing process, while engineers and technicians often did not understand each other’s work or language; and ?, who cites lack of communication as a particular obstacle to coordination between nurses and physicians in health care teams. Using an experimental investigation, ? show that when teams who develop their own codes or languages are merged, incompatibility between the common languages greatly inhibits the performance of the post-merger organization.

Against the backdrop of these motivating studies, we examine a stylized environment where success in each period (of an infinite-horizon game) depends on two agents using compatible platforms. The initial tension arises because neither party is able to discern which of the possible platform pairings are compatible. Over time, the agents may switch between different platforms in an attempt to coordinate their activities. However, there are two potential sources of friction that may hinder attempts to coordinate. First, in many settings, coordinating on compatible platforms (or procedures) may be necessary *but not sufficient* for success.¹ When the outcome is noisy, learning about compatibility is imperfect, as failing to observe success need not imply incompatibility. Second, switching between platforms is generally costly. We introduce an explicit cost of switching to represent the frictions associated with modifying procedures or transitioning between different platforms. As such, not only must the agents determine whether or not one of them needs to switch to a different platform, they must also grapple with the incentives for each to rely on the other being the one who incurs the cost of switching.

In organizational settings, these types of frictions are ever-present—for example, different organizational units may lack a common language with which to communicate their protocols and specialized knowledge to each other. Switching procedures may involve costly activities like replacing systems, purchasing new equipment, updating manuals and protocols, retraining staff, or repurposing facilities. These sources of friction within decentralized organizations can make coordination even more difficult to achieve. Accordingly, we examine three main questions

¹For example, successful innovations or scientific breakthroughs may be impossible when units employ incompatible methods. However, compatible methods do not, by themselves, always lead to breakthroughs or innovations.

that naturally apply to coordination problems, particularly within organizations. **First**, how does the presence of an explicit cost of switching impact the optimal strategies in a dynamic coordination problem? In addition to the strategic uncertainty associated with the lack of a common knowledge description of the game, each unit in the organization would prefer that the other be the one to incur the cost of switching. **Second**, how does noise in the learning process affect the optimal switching behavior and the corresponding ability of the players to achieve efficient coordination? Learning is substantially more complex if compatibility only stochastically improves the productivity of the organization. **Third**, how does the interaction between costly switching and noisy learning further hinder coordination?

To address these research questions, we examine two settings—(1) a *Deterministic* environment, in which compatible platforms guarantee success, and (2) a *Stochastic* environment, in which failure remains possible (although less likely) when the agents are on compatible platforms.² We analyze each setting theoretically and empirically, using a controlled laboratory experiment in which we vary the cost of switching and the complexity of the learning environment (Deterministic or Stochastic).

Throughout the paper, we assume players are unable to communicate. This assumption, although seemingly strong, reflects two key considerations. First, in many organizational settings, different units may not possess the requisite technical language to communicate meaningfully about their procedures. Second, communication regarding switching strategies may be infeasible or lack credibility, particularly if switching is costly. While allowing communication may (in theory) open up the possibility of asymmetric equilibrium play, it is far from obvious how parties that can communicate will practically coordinate on which asymmetric equilibrium to play when the asymmetry imposes a payoff disadvantage on one of the parties. Consequently, we abstract from communication in order to focus on how the players handle the strategic tensions of learning under the two main frictions described above. Despite this focus, there are many studies showing that communication can help to improve coordination in related strategic environments. As such, future research might explore, as a next step, whether participants can also effectively use communication to overcome the intricate strategic tensions of our environment.

Overview. In our theoretical model, two players each have two available platforms and know that there are two compatible platform pairings. Each compatible pairing generates a success (higher payoff outcome) with some known probability. Incompatible platforms guarantee failure (a lower payoff outcome). Neither party knows which configurations are compatible, and they lack a common description or language with which to understand or identify each other’s platforms. Instead, players must navigate their way to one of the compatible platform pairings over time. In every period, each player chooses simultaneously whether to switch platforms (incurring the associated cost) or remain on their current platform. At the end of the period, updated platforms are determined by the starting platforms and the players’ actions (to switch or remain), and the outcome of the organization’s project—success or failure—is publicly observed.

The strategic tensions of the problem are as follows: Initially, players are uncertain about whether or not their platforms are compatible. In the Deterministic setting, compatibility is always perfectly revealed after the first period. That is, success in the first period reveals that the players are on compatible platforms entering the second period, while failure in the

²In both settings, we assume that incompatible platforms always ensure failure.

first period reveals that they are on incompatible platforms entering the second period. In the former case—and after any observed success for that matter—maintaining coordination is straightforward since each player can simply remain on their platform in every future period. But in the latter case, the subjects face a difficult problem; coordinating their activities requires one and only one of the players to switch, but anyone who switches must incur a cost. In the event that both players switch platforms, they will end up on (different) incompatible platforms on top of both incurring the cost of switching. In a symmetric equilibrium, this tension drives the agents who are uncertain about compatibility to choose a mixed strategy, in which they switch with a lower probability as the cost of switching increases.³

In the Stochastic setting, compatibility is only revealed after success is realized. Nevertheless, players’ beliefs about the hidden state (i.e., compatibility) conditional on the updated platforms always deteriorate after the observation of a failure. In this case, whenever they are sufficiently confident that the current platforms are compatible, players should remain on their current platforms. Conversely, as beliefs deteriorate, such that players become more pessimistic that the current platforms are compatible, the symmetric equilibrium prescribes that the players choose mixed strategies, randomizing between switching and remaining on their current platforms. We show how the probability of switching in this symmetric equilibrium depends on the cost of switching, and how it increases as the belief further deteriorates.

To complement our theoretical framework, we design and conduct a controlled (virtual) laboratory experiment that implements the stylized model with treatments that vary the cost of switching and the setting (Deterministic or Stochastic). Our experimental data are broadly consistent with the comparative statics predictions derived from the theoretical model. However, there are several notable deviations from the predictions that invite more discussion. One key methodological contribution of our experiment is that we allow subjects to explicitly randomize their implemented action (facilitating mixed strategies) using a *switching rule* that corresponds to the probability of switching platforms. This allows us to observe the mixed strategies chosen by subjects, rather than having to infer them from aggregated empirical distributions, and provides a richer dataset with a cleaner mapping from the predictions to the data.

In the Deterministic setting, initial switching rules are decreasing with the cost of switching, in line with the predicted comparative statics. In contrast, in the more complex Stochastic setting, initial switching rules are only partly consistent with the comparative statics predictions. On average, switching rules in the first period are lower when switching is costly than when it is not, but do not decline when the switching cost increases from small to large.

We then provide several results concerning behavior beyond the first period of the game. First, in all treatments, the mean switching rules chosen by players once they know they are on compatible platforms are very close to, but strictly higher than, zero. That is, subjects are generally capable of playing the efficient, symmetric equilibrium strategy (both remain on their current platforms) once they are coordinated, even in the Stochastic setting where there is still some probability of failure. Nevertheless, there are a non-trivial fraction of subjects who continue to mix even when they know they are on compatible platforms. We show that these departures from the theoretical prediction are correlated with lower cognitive ability (as measured by performance on a version of the Cognitive Reflection Test). Moreover, they become substantially less frequent when subjects are more experienced. Finally, in the Stochastic

³There may also be asymmetric equilibria, although they involve otherwise symmetrically informed players arriving at the same conjecture about who should switch.

setting, we show that the departures are more likely after the parties experience a failure or a series of successive failures (by chance) than after a success. This suggests that bounded rationality or inattention may also be an important factor for maintaining coordination when outcomes are noisy.

Second, contrary to our predictions, subjects who are on “likely compatible” platforms in the costly Stochastic treatments choose positive switching rules instead of opting to remain on their current platforms. Switching rules are moderately higher when subjects are on likely incompatible platforms, but the predicted relationship between switching rules and the common belief about compatibility is attenuated relative to the theory. To explain these findings, we show that, as with behavior on known compatible platforms, the mean switching rules are especially responsive to the number of consecutive failures, regardless of whether the underlying belief is more optimistic than pessimistic.

Third, in the Deterministic treatments, subjects who learn they are on incompatible platforms choose mean switching rules that are highly consistent with both the theoretical point predictions and the comparative statics (with respect to the cost of switching). Subjects in this situation overwhelmingly take up the option to select an explicit mixed strategy and respond as predicted to an increase in the cost of switching. Conversely, when players reach known incompatibility as a result of a previous “transgression” (whereby one of the players switched their platform away from a known compatible pairing), we find some evidence that they attempt to use the asymmetric history as a precedent for future asymmetric play.

Overall, the added difficulty of learning how to coordinate in the Stochastic setting leads to slower and less frequent coordination than in the Deterministic environment. When subjects are inexperienced, this translates into higher mean payoffs (per period) in the Deterministic treatments than in their Stochastic counterparts. However, as subjects gain experience, they initiate fewer transgressions when on compatible platforms, especially in the costly Stochastic treatments. As a result, we find that the differences in mean payoffs between costly Deterministic and Stochastic treatments disappear with experience.

Relation to prior literature. Our motivation and framework bear some similarities to an approach introduced by ? to examine the problem of achieving coordination through repeated play when players lack a common language with which to describe the game. One key idea underlying their approach is that as long as players’ actions or roles cannot be distinguished based on the history of play, they ought to be constrained to use strategies that treat any such indistinguishable actions symmetrically. That is, in their framework, the strategic uncertainty generated by the lack of a common knowledge description of the game constrains the statistical relationship between the players’ strategies. As such, players must instead adopt symmetric strategies until the history of play generates asymmetric precedents upon which to distinguish some of the available actions. In this way, players can learn how to coordinate, in a decentralized manner, through repeated play.

Several other researchers have extended the main results developed by ?, including ?, ?, ?, and ?. Their results are also closely related to the operations research literature on “rendezvous search” (see, e.g., ??). Crawford and Haller model the problem by requiring players to respect all symmetries that are indistinguishable based on past play, which prescribes symmetric strategies and symmetric treatment of actions by each player, at least until some distinction emerges. Since players have perfectly aligned incentives, the problem reduces to choosing a set of strategies that

achieves coordination as quickly as possible, subject to what they refer to as an *attainability* constraint. Their formulation has been applied to other similar environments—most notably, in a “search-for-success” game studied theoretically by ? and experimentally by ?.⁴ However, to the best of our knowledge, none of these studies consider the impact of explicit switching costs (or the associated frictions they represent) on dynamic coordination.

The literature examining the costly nature of switching actions in games is relatively sparse. Moreover, the few papers that do examine switching costs concentrate on questions and behavior in settings very different from ours.⁵ Intuitively, when switching is costly, it may encourage a greater reluctance to switch platforms, even when the players are very confident that coordination requires somebody to switch. When there are natural asymmetries between the players (for instance, if one player has a lower switching cost than the other), it may be possible to use the asymmetry as a focal point (see, e.g., ???).⁶ Yet, without a common language, and without any other apparent asymmetries between the parties, it is unclear how they would synchronize their activities on a particular asymmetric equilibrium even when one exists.⁷ Thus, we concentrate in this paper on a symmetric environment, in which there are (initially) no asymmetric features to focus coordinated decisions.

One recent study related to our setting is ?. They examine an infinitely repeated volunteer’s dilemma, in which a higher payoff requires at least one party to incur a cost of volunteering. They vary the cost of volunteering in a laboratory experiment (including treatments with symmetric and asymmetric costs) and focus on the prevalence of turn-taking strategies in infinite-horizon (repeated) interactions. In the case where players have asymmetric costs of volunteering, they further examine the tension between efficiency and equity that arises in the asymmetric equilibrium prescribing that the lower cost player always volunteers. Relative to their study, our environment differs in at least two fundamental ways. First, the players in our setting must coordinate on *exactly* one player choosing to switch, and must do so only when they are on incompatible platforms. Second, the players in our infinite-horizon game must learn about compatibility over time. As such, it is often unclear whether a ‘volunteer’ (who incurs the cost of switching) is necessary at all. This contrasts with ?, in which there is no uncertainty and a volunteer is needed in every period of the game.

Outline. We organize the remainder of the paper as follows. In Section 2, we outline the stylized theoretical framework and derive equilibrium predictions for both the Deterministic and the Stochastic environments. Section 3 describes the design and procedures of the laboratory experiment and summarizes the corresponding predictions and experimental hypotheses. Our main results are reported in Section 4. We discuss additional results comparing coordination rates and payoffs between treatments in Section 5, and summarize with some concluding remarks in Section 6.

⁴Another closely related study by ? examines dynamic coordination using organizational routines when players have private information about payoffs, actions are unobservable, and communication is precluded.

⁵See the work by ????

⁶For example, the parties may coordinate on a routine whereby the one with the lower cost always switches when it is optimal for somebody to do so. Nevertheless, if costs are private and non-verifiable, it may prove difficult to implement such a routine.

⁷The difficulty of justifying such an ad-hoc meeting of the minds is at the heart of the argument made by ?, for games where players have interests that completely coincide, but also for games with private information that involve some misalignment of interests, as in ?.

2 Theoretical Framework

In this section, we present a concise description of the model and associated analysis in order to highlight the main characterizations and the predictions to be tested by our laboratory experiment. We provide a more detailed, general theoretical framework in Online Appendix A.

2.1 Setup

Consider two players (e.g., two divisions within an organization), 1 and 2, who play a dynamic game with infinite horizon. The divisions perform different functions within the organization; however, their respective activities jointly determine the success or failure of some project or task in each period. Division 1 has two available *processes*, X and Y . Likewise, Division 2 has two available *techniques*, α and β . In general, we refer to processes and techniques as *platforms* for the relevant player. As such, there are four possible platform pairings: (X, α) , (X, β) , (Y, α) , and (Y, β) .

Compatible platforms. For each player, each of their platforms is (mutually) compatible with *exactly one* of the other player’s platforms. This means that there are two possible configurations of compatible platform pairings. In one configuration, (X, α) and (Y, β) are the compatible pairings (with the other two pairings being incompatible). In the other configuration, (X, β) and (Y, α) are the compatible pairings. At the beginning of the game, the true configuration is chosen by Nature, with each of the two configurations being equally likely to be chosen. The true configuration is fixed for the entirety of the interaction. Critically, *neither player observes whether or not they are on compatible platforms at the outset of the interaction.*

When players are on incompatible platforms at the end of a period, the outcome of the project for that period is failure, with an associated benefit of $y_t = 0$ for each player.⁸ On the other hand, when players are on compatible platforms at the end of a period, they experience a success with probability $p > 0$ and failure with probability $1 - p$. We differentiate between two cases: the *Deterministic* case, in which $p = 1$, and the *Stochastic* case, in which $p < 1$. In all cases, the benefit associated with a success is identical for the two divisions, $y_t = 1$.⁹

Initial platforms. At the beginning of the game, each division is assigned to (or endowed with) an initial platform. We assume that players hold a common prior belief, μ_0 , that the initial platforms are a compatible pairing. For instance, if platforms are independently and randomly drawn with equal probability, then the common prior belief (given that players are also unaware of which platform pairings are compatible) is equal to $\mu_0 = 0.5$.

Decisions in any given period. In every period, each player chooses (possibly randomly) an action—SWITCH or REMAIN—that determines whether they switch from their current platform (e.g., process X) to the other platform (e.g., process Y). The action profile determines players’ updated platforms, and the nature of the updated platform pairing (compatible or incompatible) determines the players’ benefits for that period. In addition, any player whose action choice is SWITCH must pay a switching cost, c .

⁸More generally, the benefit from failure may be $y_L \geq 0$.

⁹Similarly to failure, this is a normalization. More generally, the benefit from success may be any $y_H > y_L$.

		?	?			?	?
Division 1's View	X	1, 1	0, 0	Division 2's View	α	1, 1	0, 0
	Y	0, 0	1, 1		β	0, 0	1, 1

Figure 1. Players' descriptions of the benefits matrix

Histories, strategies, and payoffs. At the end of each period, the players observe the realized action profile (i.e., both SWITCH, both REMAIN, only 1 chose SWITCH, or only 2 chose SWITCH) and the realized benefits, which are $(1, 1)$ when a success occurs and $(0, 0)$ when failure obtains.

At the beginning of period t , a history h_t consists of a list of all past action profiles, $(a_\tau)_{\tau=0}^{t-1}$, and the resulting benefits, $(y_\tau)_{\tau=0}^{t-1}$. For each player, a behavioral strategy specifies a probability of choosing SWITCH for each history of the infinite-horizon game. The players have a common discount factor, $\delta \in (0, 1)$, and aim to maximize the expected discounted sum of payoffs from all periods.

Although the structure of the game, as described above, is common knowledge among the players, they do not share a common language description of the game. As such, the initial uncertainty regarding platform compatibility implies that the players must learn *how* to coordinate on achieving compatible platforms. The lack of a common description of the game can be illustrated as in Figure 1 for the *Deterministic* case where $p = 1$ (that is, when compatible platforms guarantee a success). In particular, neither the other player's platforms nor the positions of the platforms in the benefits matrix are commonly described to the two players.

The player viewpoints of the benefits matrices, shown in Figure 1, are useful for highlighting the information available to each player—Division 1 knows that one of Division 2's techniques is compatible with process X, while the other is compatible with process Y; similarly, Division 2 knows that one of Division 1's processes is compatible with technique α , while the other is compatible with technique β . As depicted in Figure 1, both players view themselves as the player who chooses the row, with the other player's choice represented by the columns, which are labeled by '?' to indicate that compatible pairings are initially unknown. Thus, it would not help for the players to be able to communicate with each other that they should meet in the "top-left" cell of the benefits matrix, because (i) each of them perceives their role as choosing top or bottom, and (ii) the action described to Division 1 as "top" need not correspond to the one portrayed as "left" in Division 2's view of the game.¹⁰

2.2 Equilibrium Behavior—One-Stage Game

It is useful to begin by analyzing a one-stage version of the game. Furthermore, suppose that $p = 1$ (so that compatibility guarantees success) and let $c = 0$ (so that switching is costless). Fixing the common prior belief μ_0 , and letting s_i denote the probability that player i chooses SWITCH, we obtain the following payoffs for player i : from playing SWITCH ($a_i = S$),

$$u_i(S, s_j) = \mu_0 s_j + (1 - \mu_0)(1 - s_j) = 1 - \mu_0 - s_j(1 - 2\mu_0),$$

¹⁰For instance, given the viewpoints illustrated by Figure 1, if the players located a compatible pairing at (X, β) , Division 1 would describe it as "top-left," while Division 2 would describe it as "bottom-right," underscoring the lack of a common description of the game.

and from playing REMAIN ($a_i = R$),

$$u_i(R, s_j) = \mu_0(1 - s_j) + (1 - \mu_0)s_j = \mu_0 + s_j(1 - 2\mu_0).$$

From these payoffs, it follows that player i prefers SWITCH over REMAIN if $1 - 2\mu_0 \geq 2s_j(1 - 2\mu_0)$. When $\mu_0 > 0.5$, this condition reduces to $s_j > 0.5$; when $\mu_0 < 0.5$, it reduces to $s_j < 0.5$; whereas, when $\mu_0 = 0.5$, player i is indifferent between SWITCH and REMAIN regardless of player j 's strategy, s_j .

Therefore, in general, there are three equilibria in the one-stage version of the game (and a whole continuum when the common prior is $\mu_0 = 0.5$). When $\mu_0 > 0.5$, both (S, S) and (R, R) are Nash equilibria (in pure strategies). In addition, there is an equilibrium in non-degenerate mixed strategies, with each player mixing between S and R with equal probability. Conversely, when $\mu_0 < 0.5$, the profiles (S, R) and (R, S) are the pure strategy equilibria, with the same mixed strategy profile (equal mixing between S and R by both players) remaining an equilibrium.

How might players select among multiple equilibria? With the symmetry of the success payoffs from compatible platform pairings and the assumption that $c = 0$, there is no focal or salient equilibrium. Communication may help to facilitate coordination in these settings. However, absent communication, a compelling argument may be made that symmetric equilibria are more plausible than asymmetric equilibria. For instance, suppose that $\mu_0 < 0.5$. In this case, only the fully mixed Nash equilibrium, $s_i^* = (0.5, S; 0.5, R)$ for $i = 1, 2$, is symmetric. Thus, an appeal to symmetry is one way to resolve the equilibrium selection issue. In contrast, when $\mu_0 > 0.5$, all of the Nash equilibria are symmetric, such that the question of which equilibrium would be played persists.

Next, suppose we introduce positive switching costs, $c > 0$. For $\mu_0 \geq 0.5$, this ensures that the Nash equilibrium (R, R) is Pareto dominant.¹¹ That is, costly switching makes equilibrium selection (among symmetric equilibrium profiles) somewhat simpler. For $\mu_0 < 0.5$, it also creates asymmetry in the payoffs from the asymmetric Nash equilibria, (R, S) and (S, R) . Neither player wants to be the one to switch (and thereby incur the cost c), which arguably renders the two asymmetric pure strategy equilibria even less plausible than when $c = 0$. In fact, due to the cost of switching, there is a lower cutoff belief, equal to $(1 - c)/2$, above which (R, R) remains a Pareto dominant equilibrium. For beliefs below this cutoff, the symmetric mixed strategy Nash equilibrium involves each player choosing to SWITCH with probability strictly less than 0.5. Concretely, for $\mu_0 < (1 - c)/2$, the symmetric equilibrium switching rate is given by

$$s^*(\mu_0) = \frac{1 - c - 2\mu_0}{2(1 - 2\mu_0)}, \quad [1]$$

which is strictly decreasing in c .¹²

¹¹Note that the strategy profile (S, S) is still an equilibrium with $c > 0$. Similarly, there is also a symmetric mixed strategy equilibrium, in which each player chooses to SWITCH with probability $s_i^* < 0.5$ and s_i^* is decreasing in c . However, each of these involves inefficient costly switching and, therefore, lower payoffs.

¹²In fact, notice that if c is sufficiently large, for a given $\mu_0 < 0.5$, the asymmetric equilibria (R, S) and (S, R) will cease to exist, and the unique equilibrium will be (R, R) . That is, if switching is sufficiently costly, REMAIN becomes a dominant strategy for any belief above a certain level.

2.3 Equilibrium Behavior—Multi-Stage Dynamic Game

We now extend the analysis to the multi-stage dynamic game. In general, the relevant solution concept is subgame perfect equilibrium (SPE). However, as explained later, we predominantly focus on a subset of SPE: those in which the strategies are stationary Markov strategies for a particular class of histories—those histories for which a success has never been observed—where the players’ common belief about platform compatibility, μ_t , is the relevant state.

For the moment, continue to assume that $p = 1$. In this Deterministic case, since the outcome is fully revealing about the underlying state, the players learn immediately after the initial period whether they are on compatible or incompatible platforms. If $y_0 = 1$, then they learn that they are on compatible platforms and the players’ common updated belief is $\mu_1 = 1$ at the beginning of the next period. The Pareto dominant (efficient) equilibrium, given $c > 0$, is to proceed by playing (R, R) at all future histories at which $\mu_t = 1$ (coupled with an appropriate specification of strategies at histories for which $\mu_t = 0$).¹³ For this reason, we focus on the most efficient equilibrium, where $s_i^*(\mu_t = 1) = 0$ for both players. That is, once compatibility is achieved, players simply REMAIN with certainty. This allows for a simple calculation of the continuation payoffs for any history at which $\mu_t = 1$.

If instead the players observe $y_0 = 0$, they learn after the initial period that they are on incompatible platforms, leading to the common updated belief $\mu_1 = 0$. In this case, they know that attaining compatibility requires one and only one player to switch. But since switching is costly, each prefers the other be the one who chooses to switch. If $a_0 = (S, S)$ or $a_0 = (R, R)$, there is also no asymmetric history of play that the players might appeal to as some sort of precedent with which to correlate or coordinate their decisions. The players are in similar positions as they were prior to the initial period, albeit with a different belief ($\mu_1 = 0$ instead of μ_0), and therefore the same arguments for focusing on the symmetric equilibrium (in mixed strategies) can be invoked.

However, if initial play is asymmetric (e.g., the action profile is $a_0 = (S, R)$ or $a_0 = (R, S)$), it is theoretically possible for the players to use such asymmetric play to correlate subsequent action profiles. Consider the following argument in the case where $a_0 = (S, R)$ and $\mu_1 = 0$: “Player 1 (who chose S) *caused* the incompatibility and, therefore, the onus is on her (player 1) to switch back.” If both players subscribe to this argument, then asymmetric play may be justified. The problem persists, however, when one considers that an equally appealing argument can be made for the opposite conclusion: “Player 1 (who chose S) already incurred a switching cost and, therefore, it is only fair that player 2 switch next, since he has yet to incur any cost.” If both players subscribe to this argument, it lends some credibility to the opposite pattern of asymmetric play. Ultimately, these arguments do little to resolve the underlying strategic uncertainty that characterizes the equilibrium selection problem (absent the ability for players to communicate). Thus, the argument in favor of symmetric strategies remains compelling. Therefore, if players do not, or are unable to, condition their continuation play on the asymmetric action profile in $t = 0$, the construction of continuation payoffs (for $t = 1$) used to characterizing optimal play in $t = 0$ implies that symmetric strategies are also optimal in $t = 0$ of the dynamic game. Formally, the closed-form solution for the symmetric equilibrium

¹³For instance, one natural specification is that if $\mu_t = 0$, choose the action profile (S, R) if the most recently played asymmetric action profile is (S, R) , and play (R, S) if the most recently played asymmetric action profile is (R, S) . In other words, the player whose SWITCH action induced the incompatibility is responsible for switching back to restore compatibility.

switching rate is given by $s^*(\mu_t) = 0$ for all $\mu_t \geq \mu^{D*}$ and

$$s^*(\mu_t) = \frac{1}{2} \left(1 - \frac{c(1 - 0.5\delta)}{(1 - 2\mu_t)(1 + 0.5\delta c)} \right) \quad [2]$$

for any belief $\mu_t < \mu^{D*}$, where

$$\mu^{D*} = \frac{1}{2} \left(1 - \frac{c(1 - 0.5\delta)}{1 + 0.5\delta c} \right). \quad [3]$$

That is, given the cutoff belief μ^{D*} (which is never greater than 0.5), if the prior belief is above the cutoff or players have learned that they are on compatible platforms, the equilibrium switching rate is zero. However, when the prior is below the cutoff, or if players have learned they are on incompatible platforms, the equilibrium switching rate is strictly positive, higher for lower values of μ_t , and decreasing in the cost of switching, c .

Stochastic environment. Next, we consider how things may change if failure is not perfectly revealing of the state. Thus, we turn to the *Stochastic* case, in which $0 < p < 1$. In this case, once players observe a success ($y_t = 1$), they can perfectly infer the state in all future periods, and the same arguments for efficient equilibrium play (each player choosing REMAIN with certainty) apply. However, since $p < 1$, unless the players have previously observed a success, failure ($y_t = 0$) no longer perfectly reveals the state, since it can occur even when players are on compatible platforms. Nevertheless, since players observe the realized action profiles, beliefs about compatibility evolve in concert for the two players based on the updated belief at the end of the previous period, the action profile in the current period, and the realized outcome (success or failure). Thus, we focus on beliefs and behavior by players prior to observing a success.

The evolution of beliefs during this phase is relatively straightforward. Suppose $\mu_t \in (0, 1)$ is the common belief that players are on compatible platforms at the beginning of period t .¹⁴ If the outcome in period t is success ($y_t = 1$), the updated posterior belief is $\mu_{t+1} = 1$. However, if the outcome in period t is another failure ($y_t = 0$), the updated belief depends on the realized action profile a_t , in the following manner:

$$\mu_{t+1} = \begin{cases} \frac{\mu_t(1-p)}{1-p\mu_t}, & \text{if } a_t \in \{(S, S), (R, R)\}, \\ \frac{(1-\mu_t)(1-p)}{1-p(1-\mu_t)}, & \text{if } a_t \in \{(S, R), (R, S)\}. \end{cases}$$

Intuitively, when the players' realized actions are the same, they either remain on the same platforms or shift to the complementary pairing for the same configuration. As such, failure represents bad news about the configuration they were on to begin the period t , and the belief μ_t is revised downward to μ_{t+1} . Conversely, when the players' realized actions are misaligned, they shift to the alternative configuration of platform pairings (for which the initial belief is $1 - \mu_t$) and this belief is revised downward. Correspondingly, since the players have switched from one configuration to another, failure represents bad news about the updated platforms and therefore *good news* about the compatibility of the platforms they were on at the beginning of

¹⁴Note that this implies a success has not yet been observed.

period t .

By restricting attention to symmetric play during this phase, we effectively impose some stationary Markovian structure on the equilibrium strategies at certain histories. In particular, strategy profiles are such that, for any history where the players have not yet observed a success, the switching rule is a function of the players' common belief, μ_t , but does not otherwise depend on the history of action profiles or the period, t .¹⁵ This allows us to define continuation payoffs in terms of the common belief, μ_t . Nevertheless, there is still a small technical challenge. Since players can, in theory, experience infinitely repeating failure, there are infinitely many possible states to consider. For example, if players are on incompatible platforms, both choose to mix between SWITCH and REMAIN, and the history of action profiles is eternally (R, R) , then the belief μ_t will continue to deteriorate towards 0.

Fortunately, longer strings of successive failure become extremely unlikely. As such, continuation payoffs for beliefs that are sufficiently close to the two extremes ($\mu_t = 0$ and $\mu_t = 1$) converge to the corresponding continuation payoffs, $V(0)$ and $V(1)$. This allows us to solve recursively for an equilibrium approximation; for any $\mu_t \in (0, \varepsilon)$, we fix $V(\mu_t) = V(0)$ and for any $\mu_t \in (1 - \varepsilon, 1)$, we fix $V(\mu_t) = V(1)$, then we take $\varepsilon > 0$ to be arbitrarily small, in order to exploit the recursive structure.

Formally, a mixed Markov strategy for player i is a function $s_i : [0, 1] \rightarrow [0, 1]$ that assigns to each $\mu \in [0, 1]$ the probability with which player i chooses the action SWITCH. A profile of Markov strategies that constitute a subgame perfect equilibrium of the dynamic coordination game is called a Markov perfect equilibrium (MPE), and a symmetric MPE is one in which $s_i^*(\mu) = s_j^*(\mu)$ for all μ .

Following the arguments in favor of symmetric equilibrium play that we invoked earlier and employing the same equilibrium selection criterion as above yields a particularly intuitive characterization of (approximate) equilibrium behavior. There exists an equilibrium cutoff belief, $\mu^* \leq 0.5$, which depends on the switching cost (along with other parameters of the model). Then, for all beliefs $\mu \geq \mu^*$, both players choose to remain with certainty ($s_i^*(\mu) = 0$). For all interior beliefs below the cutoff, $0 < \mu < \mu^*$, the equilibrium switching rate $s^*(\mu)$ is strictly between 0 and 1, strictly decreasing in c , and, except when $c = 0$, strictly decreasing in μ .¹⁶ Full details of our approximation procedure are given in Online Appendix B. Figure 2 provides a visual representation of the equilibrium switching rates, as a function of μ_t , for each setting under the parameter values used in our experiment, which we describe next.

3 Experimental Design

3.1 Setup and Treatments

In our experiment, subjects played a series of supergames—each one an indefinitely repeated two-player game—designed to reproduce the theoretical framework described above. The length of each supergame was determined according to a random termination rule, with continuation probability $\delta = 0.9$.¹⁷ In each supergame, prior to the first period, subjects were randomly

¹⁵Note that we need not make the same restriction on strategies following an observed success. For instance, if players reach the belief $\mu_t = 0$, which can only occur after having observed a success and subsequently shifted to incompatible platforms, we may wish to allow for play to be conditioned on the identity of the player whose decision to switch resulted in the incompatibility.

¹⁶When $c = 0$, the equilibrium switching rate is $s^*(\mu) = 0.5$ for all $\mu \in (0, \mu^*)$.

¹⁷Further details on the random termination procedure are described in Section 3.2.

divided into pairs.

Each player in a pair was assigned to an initial platform, represented by a color. For one player, the available color platforms were green and purple, while for the other player, the available platforms were blue and orange. At the beginning of each supergame, platforms were matched (one-to-one) into compatible pairings—either (i) (green, blue) and (purple, orange) or (ii) (green, orange) and (purple, blue). Players were not informed about which platform pairings were compatible, only that each configuration was equally likely. Furthermore, neither player was ever informed about the other player’s color platform. The probability that players were initially assigned compatible platforms was exogenously fixed to $\mu_0 = 1/3$, and this was common knowledge among the players.

In each period, players could choose between two available actions, SWITCH and REMAIN, which determined whether their platform would switch or stay the same color it was at the beginning of the period. The cost to choosing REMAIN was zero, whereas a player who chose SWITCH incurred a cost equal to c , which we varied across treatments. In addition to the two pure actions, we also allowed players to *explicitly* randomize their action. On the decision screen, subjects were given an entry box in which they could indicate a *switching rule*, which could be zero (corresponding to REMAIN), 100 (corresponding to SWITCH), or any integer in between. The subject’s action was then determined by a random integer, drawn uniformly from between 1 and 100 (inclusive)—if the random integer was equal to or below her selected switching rule, then her action was SWITCH and she incurred the cost c ; otherwise, her action was REMAIN and no cost was incurred.¹⁸ Thus, the randomization procedure allowed a subject to choose, via the switching rule, the probability with which she wanted her action to be SWITCH.

At the end of the period, each player’s payoff was determined by two things: (i) the benefit generated by the players’ updated platforms and (ii) the cost of their implemented action. When the updated platforms were incompatible, each player received a low benefit of $y^L = 120$ points. When the updated platforms were compatible, the benefits were y^H points for each player with probability p and y^L points each with probability $1 - p$. The values of y^H and p were varied across treatments (see below).

Each player’s payoff was determined by subtracting the cost of her implemented action from her realized benefit. Thus, if the player’s action was SWITCH, her payoff was equal to her benefit minus the cost, c . If the player’s action was REMAIN, her payoff was simply equal to her benefit. Total payoffs from a supergame were calculated by summing together the payoffs from all periods before the supergame terminated.

Treatments. In one set of treatments (*Deterministic*), we set $y^H = 190$ points and $p = 1$, so that compatible platforms always generated the high benefit. In the other set of treatments (*Stochastic*), we set $y^H = 260$ points and $p = 0.5$. To maximize comparability, these parameters (together with $y^L = 120$ points) ensured that in all of our treatments, the expected benefit for players on compatible platforms was equal to 190 points. Within each set of treatments, we considered three different values for the cost of SWITCH: $c = 0$, $c = 10$, and $c = 30$ points. Thus, we examined six treatments in a 2×3 design (see Table 1), referred to as D0, D10, D30

¹⁸To help subjects visualize and select a switching rule, the decision screen displayed a 10×10 grid of boxes labeled with all integer numbers from 1 to 100. After a subject entered a switching rule, the boxes with numbers equal to or below the switching rule were highlighted in yellow, while those with numbers greater than the switching rule were shaded white. Subjects could also click the SWITCH or REMAIN buttons to directly select one of the two pure actions to be implemented. Screenshots are provided in Online Appendix F.

Table 1. Summary of experimental treatments.

Treatment	Setting	c	Sessions	Subjects	Groups
D0	<i>Deterministic</i>	0	2	28	4
D10	<i>Deterministic</i>	10	4	48	8
D30	<i>Deterministic</i>	30	4	54	8
S0	<i>Stochastic</i>	0	2	28	4
S10	<i>Stochastic</i>	10	4	56	8
S30	<i>Stochastic</i>	30	4	52	7
Total			20	266	39

and S0, S10, S30. We collected data from two sessions each in the D0 and S0 treatments, and from four sessions each in the other four treatments. The number of completely independent matching groups (shown in the far right column) is small, which may cause some concerns about statistical power. Nevertheless, these are very conservatively defined—they do not represent the number of pairs in each treatment, all of which are considerably larger, but the matching groups from which the pairs are drawn. While there may be some dependence across pairs drawn from the same matching group, this should help to alleviate some of the potential limitations due to the smaller number of observations.

3.2 Session Details and Procedures

The experiment was implemented in z-Tree (?), and all sessions were conducted virtually using z-Tree unleashed (?). For each session, we recruited between 10 and 16 subjects via ORSEE (?) from a sub-population of pre-registered students at Florida State University.¹⁹ Subjects were checked in to a Zoom meeting one at a time and assigned a numerical ID (to which they were renamed) so as to preserve anonymity. Participants could only chat with the experimenters during the session. Instructions were displayed in stages on the screen (through z-Tree) with the same experimenter reading aloud from a script to accompany each stage.²⁰ Screenshots of the instructions (with the associated script) are provided in Online Appendix F.

At the beginning of the session, the program randomly divided subjects into independent matching groups.²¹ All interactions were confined to be within a matching group. Subjects were randomly paired (within their matching group) for the first of a series of supergames (referred to as “matches”) and were randomly rematched into pairs for each subsequent supergame, following a strangers matching protocol. Although we planned for up to five supergames in each session, the two-hour time constraint on sessions prevented the subjects from reaching the fifth supergame in half of the sessions. Nevertheless, subjects in every session completed at least four supergames.

¹⁹Although we invited at least 16 subjects for every session, the number of participants who actually showed up to participate varied across sessions. We ran one session with 10 participants, ten sessions with 12 participants, four sessions with 14 participants, and five sessions with 16 participants.

²⁰The instructions included partial (interactive) screenshots and provided subjects with the opportunity to familiarize themselves with the interface. There were also two separate comprehension quizzes – one regarding the calculation of payoffs and one regarding the timing and length of a supergame.

²¹In the session with 10 subjects, there was only one matching group. In sessions with 12 or 16 subjects, there were two equally sized matching groups, while in sessions with 14 subjects, there was one matching group of 6 and one matching group of 8.

Table 2. Supergame lengths.

Sequence	Supergame					Total
	1	2	3	4	5	
1	10 (10)	15 (15)	7 (10)	16 (16)	5 (10)	53 (61)
2	9 (10)	10 (10)	16 (16)	5 (10)	9 (10)	49 (56)

Notes: The values in the table are the actual supergame lengths, with the corresponding number of decision periods reported in parentheses.

Supergame lengths & random termination. Prior to running any sessions, we generated two random sequences of supergame lengths using the continuation probability of $\delta = 0.9$. The sequences are reported in Table 2. These sequences determined the actual lengths of each supergame.²² However, we employed a variant of the Block Random Termination (BRT) rule for the *decision periods* in each supergame (see, e.g., ?). Using the BRT procedure, subjects always made decisions for a fixed block of 10 periods, even when the actual number of periods in the supergame was fewer than 10. At the end of the 10-period block, subjects were informed whether the supergame had ended during the block and, if so, after which period. If the supergame did not end during the block, subjects played another period and, after each additional period, were informed whether the supergame ended. To make this clear, the number of decision periods for each supergame is reported in parentheses in Table 2. A particular advantage of the BRT procedure, and our primary motivation for using it, is that it allowed us to observe *at least* 10 decisions for every supergame, even if the supergame ended after only a few periods.

At the end of each period, players were shown a summary screen that reported the implemented action profile (e.g., (SWITCH, SWITCH) or (SWITCH, REMAIN)), the realized benefit from their updated platforms, and their own payoff for the period. In addition, players were reminded that, if the paid periods had not already terminated, the probability that the match would continue for at least one more period was $\delta = 0.9$. After Period 10, subjects were informed whether the supergame had already ended and, if so, after which period. Moreover, when subjects were informed that the supergame had ended, they were also shown a summary of their payoffs from each period and their total payoff from the supergame before proceeding to the next supergame.

We conducted four sessions for each of the treatments with a strictly positive cost (D10, D30, S10, and S30). Two sessions were conducted using Sequence 1 of supergame lengths, while the other two sessions used Sequence 2 (cf. Table 2). For the $c = 0$ treatments (D0 and S0), we conducted one session with each sequence (two sessions total for each treatment). Thus, we conducted 20 sessions, with a total of 266 participants.

Additional tasks. In addition to the main part of the experiment described above, we also elicited two measures of individuals' characteristics: social preference, using the Allocation Game of ?, and cognitive process, using a selection of Cognitive Reflection Test (CRT) ques-

²²For each sequence, we drew a series of random integers between 1 and 100 using the website random.org. The length of the first supergame was determined by counting the number of draws up until the first draw over 90, the second supergame length by counting the number of draws until the next draw over 90, and so on. The series of random numbers drawn for each sequence is available from the authors upon request.

tions.²³ Details about the Allocation Game and CRT questions are provided in Online Appendix E. We also collected basic demographic information in the post-experiment questionnaire.

Payment. At the end of the experiment, one supergame was drawn at random for each subject and used to determine the subject’s payment. Payoffs were converted from “points” into US dollars according to the exchange rate 180 points = \$1. In addition, subjects were paid for the incentivized measures of individual characteristics and received a participation fee equal to \$10. Average earnings (including the participation fee) were \$23.80, ranging from a low of \$13.40 to a high of \$36.70.

3.3 Predictions

Derived from our theoretical analysis of the Deterministic and Stochastic cases, Figure 2 shows point predictions for the equilibrium switching rule by (common) belief μ_t , across our six experimental treatments.²⁴ For all treatments, we maintain the assumption that subjects select the most *efficient* equilibrium strategy at beliefs $\mu_t \geq 0.5$. Thus, the predicted switching rule for all common beliefs from $\mu_t = 0.5$ up to $\mu_t = 1$ is $s(\mu_t) = 0$.

Recall that in the first period, the common prior belief is exogenously imposed to be $\mu_0 = 1/3$. In the treatments with $c = 0$, the equilibrium prediction is $s(\mu_0) = 50$, just as it is for all lower beliefs. For the other treatments, the predicted equilibrium switching rule at $\mu_0 = 1/3$ is lower and decreasing in the cost of switching—from $s = 38.9$ in D10 to $s = 20.4$ in D30 and from $s = 34.9$ in S10 to $s = 11.9$ in S30. Thus, in each case, the cost is sufficiently low that the equilibrium cutoff belief is between the prior μ_0 and 0.5.

For the Stochastic treatments, the values of $\mu_0 = 1/3$ and $p = 0.5$ generate a particular set of feasible beliefs; aside from the beliefs that are within $\varepsilon = 0.05$ of the two extremes, the players’ common belief μ_t transitions between the following set of values (and their complements): $\mu_t = 1/9$, $\mu_t = 1/5$, $\mu_t = 1/3$, and $\mu_t = 1/2$. As described in the theoretical framework above, the equilibrium switching rule is increasing as the belief μ_t deteriorates towards 0. Based on these predictions, we form the following experimental hypotheses.

Hypothesis 1. *Initial switching rules at $\mu_0 = 1/3$ (in Period 1 of a supergame) are decreasing in switching cost for both the Deterministic and Stochastic treatments.*

Hypothesis 2. *In all treatments, subjects choose $s = 0$ (REMAIN with certainty) if their common belief that the platforms are compatible is $\mu_t \geq 0.5$.*

For both sets of treatments, Hypothesis 2 entails the prediction that when they know they are on compatible platforms ($\mu_t = 1$), subjects choose $s = 0$. In addition, for the *Stochastic* treatments, Hypothesis 2 implies that subjects also choose $s = 0$ when they believe their platforms are likely compatible ($0.5 \leq \mu_t < 1$).

In contrast, for the S10 and S30 treatments, the equilibrium switching rule for beliefs below $\mu_t = 0.5$ increases (non-linearly) as the belief deteriorates towards zero.

²³With the exception of one S10 session, these tasks were completed at the beginning of the session, before the main part was introduced. In the one exception, they were completed after the main part. Nevertheless, in all sessions, feedback and payoffs from the Allocation Game and the CRT questions were withheld until the end of the session.

²⁴Rather than plot the equilibrium probability of switching, between 0 and 1, we plot the corresponding equilibrium switching rule for subjects, who chose their switching rule from the set of integers between 0 and 100. Both graphs were produced using the parameter values from our experiment. For the Stochastic setting, we used a belief threshold of $\underline{\mu} = 0.05$ to compute the equilibrium approximation (see Online Appendix B for details).




Figure 2. Predicted equilibrium switching rules by (common) belief, μ_t , for parameters used in our experiment. Prior belief: $\mu_0 = 0.333$.

Hypothesis 3. *In S10 and S30, subjects choose progressively higher switching rules as the common belief μ_t deteriorates towards zero.*

Finally, when the players know they are on incompatible platforms ($\mu_t = 0$), the equilibrium switching rules are quite close together at $s = 50$ for D0 and S0, $s = 46.3$ for D10 and S10, and $s = 40.1$ for D30 and S30.²⁵

Hypothesis 4. *In both the Deterministic and Stochastic treatments, the cost of switching has a weaker, although still negative, effect on switching rules when players know they are on incompatible platforms ($\mu_t = 0$) than at the prior belief ($\mu_0 = 1/3$).*

4 Results

In this section, we present our main experimental findings. We begin by examining the initial switching rules chosen by players in the first period of the interaction. Then, we examine behavior when players are on known compatible platforms (Section 4.2), when players hold intermediate beliefs in the Stochastic treatments (Section 4.3), and when they are on known incompatible platforms (Section 4.4).

4.1 Initial Switching Rules (Period 1)

In the first period of each supergame, subjects face the common prior belief, $\mu_0 = 1/3$, regarding compatibility. According to Hypothesis 1, we expect the initial switching rules to decrease in the switching cost for both the Deterministic and Stochastic treatments. Our first main result

²⁵However, note that $\mu_t = 0$ can only be reached off the equilibrium path in the Stochastic treatments. Thus, while the symmetric predictions at $\mu_t = 0$ coincide with the symmetric predictions for the corresponding Deterministic treatments, we anticipate more instances of asymmetric play at $\mu_t = 0$ in the Stochastic treatments, where players can potentially condition their behavior on the identity of the player whose decision to switch led to their subsequent incompatibility.

is that while we cannot reject Hypothesis 1 for the Deterministic treatments, it is refuted by the behavior observed in our Stochastic treatments.

Result 1. *Initial switching rules are decreasing in switching cost for the Deterministic treatments. In the Stochastic treatments, although initial switching rules are lower in S10 and S30 than in S0, they are not lower in S30 than in S10.*

We provide support for this result below, focusing first on the Deterministic treatments, then on the Stochastic treatments.


Deterministic treatments. Figure 3a summarizes the mean switching rule in Period 1 for each of the three Deterministic treatments. The darker (navy) bars use only data from Match 1, while the lighter (gray) bars use data from all four matches. Although the mean switching rule is below the equilibrium prediction for all treatments, the comparative statics predictions are well supported by the data. Thus, we are unable to reject Hypothesis 1, whether we use all four matches or only Match 1. These main treatment comparisons are also supported by a series of Mann-Whitney Wilcoxon (MWW) Ranksum tests, using each subject’s first period (Match 1 only) switching rule as a single observation. The tests indicate significantly higher switching rules in D0 than in D30 ($p < 0.001$) and in D10 than in D30 ($p < 0.01$), although the difference between D0 and D10 is not statistically significant ($p = 0.159$) using this test.

To give a better sense of the distribution of initial switching rules, we also provide empirical CDFs of the mean initial switching rule (computed for each subject by averaging across all four matches) in Figure 3b. Consistent with the comparative statics result, the distributions can be clearly ordered according to first-order stochastic dominance.²⁶ Moreover, the empirical CDFs suggest that the differences between treatments can also be attributed in part to variation in the fraction of subjects who choose $s = 0$ in all four matches. In D30, where switching is most expensive, this is the chosen strategy for 31 out of 54 subjects (57.4%). In D10, the fraction decreases substantially, to 15 out of 48 subjects (31.3%), while for D0, where switching is costless, the fraction decreases even further, with only 3 out of 28 (10.7%) choosing $s = 0$ in all four matches.


Stochastic treatments. Figure 4a shows that the mean switching rules in Period 1 are well below the theoretical point predictions for the S0 and S10 treatments. In contrast, the mean initial switching rule in S30 is much higher than the theoretical prediction when we restrict attention to Match 1 only, and even remains slightly above the prediction when we average across all four matches. In fact, while the mean switching rule is higher for S0 than for S10, the prediction that the switching rule should be higher in S10 than in S30 is clearly not supported. Regarding the comparative statics predictions, we report a series of MWW Ranksum tests, using each subject’s first period (Match 1 only) switching rule as a single observation. Consistent with Figure 4a, the Ranksum tests indicate significantly higher switching rules in S0 than in S10 ($p = 0.034$) and no significant differences between S10 and S30 ($p = 0.699$). However, the apparent differences between switching rules in S0 and S30 are not statistically significant ($p = 0.106$).

This finding is further supported by Figure 4b, which compares the empirical CDFs of the mean initial switching rule for each subject across all four matches. While the CDFs for S10

²⁶Some caution should be exercised when interpreting these figures, since behavior in later matches may be influenced by an individual subject’s experiences in earlier matches.



(a) Mean initial switching rule




(b) Empirical CDFs


Figure 3. Mean switching rule in Period 1, by treatment (Deterministic only).

Notes: Navy bars use Match 1 only, while light gray bars use Matches 1–4. Error bars indicate standard error of the mean. Dashed (red) lines indicate equilibrium point predictions. Means for the empirical CDFs are computed at the subject level, averaging across all four matches.

and S30 are each first-order stochastically dominated by S0, the former are not significantly different from one another. In fact, a relatively higher fraction of subjects choose $s = 0$ in all four matches for S10 (57.1%, 32 out of 56) than for S30 (51.9%, 27 out of 52).



(a) Mean initial switching rule



(b) Empirical CDFs

Figure 4. Mean switching rule in Period 1, by treatment (Stochastic only).

Notes: Navy bars use Match 1 only, while light gray bars use Matches 1–4. Error bars indicate standard error of the mean. Dashed (red) lines indicate equilibrium point predictions. Means for the empirical CDFs are computed at the subject level, averaging across all four matches.

Discussion. The analysis of initial switching rules raises two particularly interesting patterns of behavior. First, the switching rules tend to be lower, on average, than the theoretical point predictions, except in the S30 treatment, where the prediction is very low. Second, relative to the patterns observed in the Deterministic setting, the main differences in the Stochastic setting appear to be driven by the S10 treatment. In particular, the fraction of subjects who choose

$s = 0$ in all four matches is much higher in S10 than in D10, and is even higher than in S30. How can we explain these two patterns of behavior? We consider three alternative explanations.

(i) Misspecified prior beliefs. One potential explanation for the relatively lower switching rules is that participants may hold initial subjective beliefs that are more optimistic than the (induced) prior explained to them in the experimental instructions. Although we explained to subjects that the prior probability of being on compatible platforms is $1/3$, it’s possible that some of them failed to incorporate this information—perhaps due to explicit optimism or inattention—and instead adopted a misspecified 50-50 prior belief. Given that the theoretical prediction for such a belief, $\mu = 0.5$, is to choose REMAIN with certainty, this would bias the mean initial switching rule towards zero and explain some of the decisions to choose $s = 0$ in all four matches. However, it’s unclear why the disparity between subjective and induced priors would be more likely in S10 than in the other treatments. Thus, while this explanation might help to explain the tendency for switching rules to be below the theoretical point predictions, it does not, on its own, explain why the disparity is so much wider in S10.

(ii) Complexity of costly and noisy learning. A second potential explanation is that, relative to the Deterministic setting, the complexity of costly learning is higher in the Stochastic setting. Participants in the Deterministic setting face a far simpler initial decision, as they learn immediately after the first period. Furthermore, the attention demands are quite minimal, as the players can easily track the compatibility of their platforms in every subsequent period.

By contrast, in the Stochastic treatments, posterior beliefs are more likely to remain interior after the first period, generating persistent uncertainty about the compatibility of their platforms. This reduces the immediate “value” of switching as the learning process is muted by the noise of the stochastic outcome. While this is already reflected in the theoretical predictions, there may be additional behavioral impacts of this added complexity if subjects need to devote more attention to the problem or are otherwise averse to persistent ambiguity (regarding compatibility). While we view these as compelling explanations for why switching rules are more conservative in the S10 treatment, it should be noted that our experiment is not explicitly designed to test for these effects.

(iii) Implicit self-randomization. Finally, although we provided participants with the option to use the explicit randomization mechanism of our experimental interface, it is always possible for subjects to perform their own internal randomization. As such, it is not possible to perfectly distinguish between a subject who intends to play the pure action, REMAIN, and one who follows a mixed strategy using internal randomization and for whom, by chance, the action REMAIN is always prescribed. With this in mind, the mean switching rules measured using our experimental data serve as a lower bound on the true mean switching rule.

To summarize, the disparities between mean switching rules and the theoretical predictions of the model may be attributed to multiple factors. Although our experiment is not designed to distinguish between these different possible explanations, we view the effects of the added complexity of costly and noisy learning, perhaps together with the effects of inattention, as the most compelling.




Figure 5. Mean switching rule for players on known compatible ($\mu_t = 1$) platforms.

Notes: Error bars indicate standard error of the mean. Dashed (red) lines at $s = 0$ indicate equilibrium point predictions.

4.2 Known Compatible Platforms ($\mu_t = 1$)

Next, we consider behavior by subjects who know they are on compatible platforms, meaning $\mu_t = 1$. Recall the prediction that, in all treatments, subjects choose REMAIN with certainty ($s = 0$) whenever they are on known compatible platforms (cf. Hypothesis 2).

Figure 5 shows that the mean switching rules chosen by players on known compatible platforms are indeed very close to the prediction in both the Deterministic treatments (left panel) and Stochastic treatments (right panel). This is particularly evident for subjects in the treatments where switching is most costly (D30, mean = 3.16 and S30, mean = 5.15). Furthermore, the percentage of subjects who chose $s = 0$ *every single time* they were on known compatible platforms (across all four matches) is substantial—46.4% in D0 (13 out of 28), 41.7% in D10 (20 out of 48), 68.5% in D30 (37 out of 54), 35.7% in S0 (10 out of 28), 46.4% in S10 (26 out of 56), and 61.5% in S30 (32 out of 52).²⁷

While the statistics listed above are based on all observations where $\mu_t = 1$, we also examine the switching rules chosen by subjects who are on known compatible platforms and have never previously shifted from known compatibility back to known incompatibility.²⁸ We refer to such situations as previous “transgressions.” When we restrict attention to observations for which no previous transgressions have occurred, behavior is even more closely aligned with the prediction. In the Deterministic treatments, the mean switching rules fall to 1.84 in D30, 5.35 in D10, and 5.21 in D0, with 72.2% (D30), 47.9% (D10), and 57.1% (D0) of subjects, respectively, choosing $s = 0$ in every possible instance. In the Stochastic treatments, the same restriction (to periods without a previous transgression) implies mean switching rules of 2.91 in S30, 5.64 in S10, and 12.96 in S0, while the percentages of subjects choosing $s = 0$ in every instance increase to 50.0% in S10 and 67.3% in S30.

²⁷Empirical CDFs for the mean (subject-level) switching rule at $\mu_t = 1$ are provided in Online Appendix D. In the Deterministic treatments, the distributions for D0 and D10 are similar, and each first-order stochastically dominates the distribution for D30. Meanwhile, in the Stochastic treatments, the distributions are all distinct—S0 exhibits FOSD with respect to S10, which in turn exhibits FOSD with respect to S30.

²⁸Experiencing known incompatibility in the past due to one of the players switching away from (known) compatible platforms may introduce a heightened degree of strategic uncertainty (doubts or concerns about the other player mixing) that leads subjects to deviate from $s = 0$ in later instances of known compatibility.

Although behavior is very close to the prediction, it is not exactly zero and significant fractions of subjects in each treatment deviate from choosing REMAIN with certainty when they are on known compatible platforms, contrary to Hypothesis 2.

Result 2. *In all treatments, the mean switching rules chosen by subjects who are on **known compatible** platforms (where $\mu_t = 1$) are often greater than zero, contrary to Hypothesis 2. This holds even when we exclude observations where subjects had previously experienced a transgression that shifted them away from known compatible platforms.*

Discussion. There may be several reasons for these departures from the predictions. One possibility is that they are the product of mistakes or the result of subjects being limited in their cognitive abilities. To explore this idea further, we compare the mean switching rules for subjects who provided more correct responses on the Cognitive Reflection Test (CRT) administered at the beginning of each session. Indeed, we find that the mean switching rules are systematically higher for those who scored lower on the CRT questions (and likewise, higher for those who provided more *impulsive* responses). The results are also robust to the exclusion of observations that occur after a previous transgression. Full details of our analysis examining switching rules and CRT responses are relegated to Online Appendix C.

Another possibility is that some subjects, faced with strategic uncertainty, play according to the (inefficient) alternative equilibrium that prescribes mutual switching when on known compatible platforms.²⁹ However, to the extent that this behavior arises, we might expect it to disappear as the subjects gain experience. Accordingly, to investigate whether mistakes or strategic uncertainty are alleviated with experience, we restrict attention to observations in Match 4 only, and to periods of known compatibility that occur prior to any transgression. Compared with the fractions reported above, the percentage of subjects who always choose $s = 0$ when $\mu_t = 1$ is substantially higher—increasing to 82.8% in D0, 88.7% in D10, 95.9% in D30, 93.1% in S0, 84.5% in S10, and 96.7% in S30.³⁰ Thus, by the time the subjects reach Match 4, switching rules by players when they are on known compatible platforms (before any transgressions) are significantly more in line with our theoretical predictions.

Finally, in the Stochastic treatments, it is also possible for players who are on known compatible platforms to experience failures. In light of this, we also study whether players on known compatible platforms choose higher switching rules after experiencing a failure or a string of consecutive failures. This may occur if players suffer from a form of recency bias, or are inattentive to the underlying compatibility, responding only to the realized payoff outcome, rather than the (known) underlying state. Consistent with this argument, we find that the mean switching rules are higher after a failure than after a success in all three Stochastic treatments (Figure 6), and also tend to increase with the number of consecutive failures, especially in the costly treatments, S10 and S30 (see Figure 7). Thus, we find some evidence that in the Stochastic treatment, repeated failures might be driving subjects to increase their switching rules, leading to a greater chance that they inadvertently switch to incompatible platforms.

²⁹We argued in Section 2 that equilibrium selection favors both players choosing REMAIN, especially when $c > 0$, as it is Pareto superior.

³⁰An additional 7.7% have a mean switching rule of just 1 (out of 100) in S10.




Figure 6. Mean switching rules by the previous outcome (success vs. failure) while on known compatible platforms ($\mu_t = 1$) in the Stochastic treatments.

Notes: Error bars indicate standard error of the mean. Dashed (red) lines indicate equilibrium point predictions. The figure excludes observations from periods after a previous transgression.




Figure 7. Mean switching rules by the number of consecutive failures while on known compatible platforms ($\mu_t = 1$) in the Stochastic treatments.

Notes: Error bars indicate standard error of the mean. Dashed (red) lines indicate equilibrium point predictions. The figure excludes observations from periods after a previous transgression.

4.3 Intermediate Beliefs in the Stochastic Treatments ($0 < \mu_t < 1$)

In this section, we examine switching rules by players in the Stochastic treatments, when they are at intermediate beliefs corresponding to “likely compatible” platforms ($0.5 \leq \mu_t < 1$) and “likely incompatible” platforms ($0 < \mu_t < 0.5$). In the former case, our theoretical predictions (Hypothesis 2) suggest that, as when they are on known compatible platforms, players should choose $s = 0$ (i.e., REMAIN with certainty). In the latter case, when players hold pessimistic beliefs, the predicted switching rule is strictly positive, increasing as the common belief deteriorates towards zero (Hypothesis 3), and higher (for any fixed belief) when the cost of switching




Figure 8. Mean switching rule by (common) belief, μ_t , in S10 (dark bars) and S30 (light bars).

Notes: Error bars indicate standard error of the mean. Dashed (red) lines indicate equilibrium point predictions.

is lower. Throughout this section, we focus on the two costly treatments, S10 and S30.³¹

In Figure 8, we plot the mean switching rule against the players' beliefs, μ_t , for S10 and S30. Given the small number of observations for beliefs near the extremes, we pool together all observations sufficiently close to zero, $\mu_t \in (0, 0.1)$ and, similarly, pool together all observations sufficiently close to one, $\mu_t \in (0.9, 1)$.³² Although the mean switching rule is generally higher at beliefs closer to zero, and lower at beliefs closer to one, the mean switching rules are not very sensitive to beliefs in the middle. For beliefs below 0.5, the mean switching rules are lower than the theoretical point predictions, while for beliefs at 0.5 or above, subjects choose higher mean switching rules than the theoretical prediction of $s = 0$.

Result 3. *Contrary to Hypothesis 2, the mean switching rules by players on likely compatible platforms are significantly higher than zero.*

Although this seems discouraging, there are many subjects whose behavior is consistent with the predictions. In particular, the fraction of subjects who *always* choose $s = 0$ when they are on likely compatible platforms is greater than 50% for both S10 and S30, as illustrated by the empirical CDFs in Figure 9. By contrast, the fraction of subjects who do so when they are on likely *incompatible* platforms is less than 20% in both treatments. Using a Sign Rank test to compare the (subject-level) mean switching rules indicates that they are significantly higher for beliefs that correspond to “likely incompatible” than for beliefs corresponding to “likely compatible” (Sign rank test, S10: $n = 55$, $Z = 3.423$, $p = 0.0006$, and S30: $n = 47$, $Z = 2.205$, $p = 0.0274$).

³¹Corresponding figures related to the analysis for S0 are provided in Online Appendix D.

³²In theory, there are infinitely many possible values of μ_t , since a string of successive failures with certain action profiles can drive the belief arbitrarily close to the extreme values without ever reaching them. However, the likelihood of such sequences is negligible and, more importantly, there are precise bounds implied by the length of the supergames played by our subjects.




Figure 9. Empirical CDFs for the (subject-level) mean switching rules chosen by players on (pooled) likely incompatible versus (pooled) likely compatible platforms, in S10 and S30.

Focusing on behavior at the likely incompatible beliefs, we do not find much support for Hypothesis 3. According to Figure 8, the mean switching rules in S10 and S30 are weakly increasing as the beliefs deteriorate towards zero, but the increase is quite modest and the means are far away from the theoretical point predictions.

Result 4. *In S10 and S30, the mean switching rules are increasing as the belief deteriorates towards zero, but the increases are smaller and more gradual than predicted by the theory.*

Thus, while behavior moves in the direction predicted by Hypothesis 3, the effect of deteriorating beliefs on mean switching rules is muted relative to theory. Considering this result together with Result 3, we turn our attention to potential explanations.

Discussion. To explain these findings, we consider the potential impacts of bounded rationality or inattention on how subjects update beliefs when outcomes are noisy. There is ample empirical evidence that individuals often struggle to compute Bayesian posteriors, particularly in complex and dynamic environments (see, e.g., ?????????). Since we did not elicit beliefs about compatibility in the experiment, we cannot test directly for systematic belief distortions.

However, we can alternatively examine whether switching behavior is more sensitive to the number of consecutive failures than to the underlying common belief. If players are more focused on the outcome (repeated failure) and inattentive towards the updated platforms, it could explain some of the switching that occurs at “likely compatible” beliefs and provide a rationale for the more attenuated response to beliefs.

In Figures 10 and 11, we plot the mean switching rule against the number of consecutive failures, separated by whether the beliefs are likely incompatible or likely compatible, for the S10 and S30 treatments, respectively. As should be expected, the mean switching rules are increasing




Figure 10. Mean switching rule by number of consecutive failures in S10.

Notes: Error bars indicate standard error of the mean.



Figure 11. Mean switching rule by number of consecutive failures in S30.

Notes: Error bars indicate standard error of the mean.

in the number of consecutive failures when we focus on the likely incompatible beliefs ($\mu < 0.5$). This is evident for both the S10 and S30 treatments—once subjects have experienced at least two consecutive failures, their mean switching rules increase significantly relative to those after experiencing zero or one previous failure.³³

In the S10 treatment, there is a more modest increase in mean switching rules with the number of consecutive failures when subjects are at likely compatible beliefs. By contrast, in the S30 treatment, there is a large increase in the mean switching rule when the number of consecutive failures increases from one to two, but the mean drops off significantly thereafter.

³³For very long sequences of failures, including those exceeding nine, there are too few observations to make meaningful comparisons.

These patterns lend credence to the argument that the flatter relationship between mean switching rules and the (implied) intermediate belief may be caused by inattention or bounded rationality, whereby subjects fail to incorporate all of the noisy information about updated platforms, and focus mainly on the growing number of failures. Similarly, they help to rationalize the higher switching rules when beliefs fall into the likely compatible range, $\mu \geq 0.5$, echoing the effect of stochastic failures on behavior when subjects are on *known* compatible platforms.

4.4 Known Incompatible Platforms ($\mu_t = 0$)

In this section, we examine the comparative statics predictions for switching rules when players know they are on incompatible platforms, i.e., where $\mu_t = 0$. Moreover, we separate the analysis into two cases: (i) those observations where $\mu_t = 0$ and the players have not previously observed a success, and (ii) observations where $\mu_t = 0$ follows a past realization of success. In the Deterministic setting, both cases (i) and (ii) are possible, whereas in the Stochastic setting, observations with $\mu_t = 0$ can only arise under case (ii).

We compare behavior in case (i) with the predictions in Hypothesis 4. In contrast, case (ii) requires that the players were previously on compatible platforms, observed a success, and subsequently shifted to incompatible platforms. This implies that one of the players must have chosen SWITCH (while the other chose REMAIN) at $\mu = 1$, triggering the subsequent incompatibility. As in the discussion of Result 2 above, we refer to this as a “transgression” and separately label players as the “transgressor” or the “compliant player.” This further allows us to examine whether the players in a given pair use the past transgression as a precedent on which to coordinate asymmetric play.

Case (i): Prior to any success. In the Deterministic treatments, behavior at $\mu_t = 0$ when subjects have not yet observed a success is highly consistent with the theoretical predictions that guide Hypothesis 4. We plot the mean switching rule by treatment in Figure 12a, which shows that, in line with the comparative statics predictions, switching rules are decreasing with the switching cost.³⁴ In addition, Figure 12b highlights the differences between empirical CDFs of the (subject-level) mean switching rules for the three cost levels. The CDFs also provide evidence that the vast majority of subjects choose (on average) interior switching rules when confronted with known incompatible platforms, as the percentage of subjects with a mean switching rule of zero is less than 20% for all three treatments.

Result 5. *In the Deterministic treatments, players who are on known incompatible platforms prior to any success choose mean switching rules that are decreasing with the cost of switching. As such, we cannot reject Hypothesis 4 for the Deterministic treatments.*

It is worth emphasizing the observation that subjects overwhelmingly use non-degenerate mixed strategies in this scenario. From a methodological standpoint, our experiment shows that eliciting mixed strategies, even for difficult coordination games, is possible. Our methodology builds on the method for eliciting randomized actions in ?, who study strategies in an indefinitely repeated Prisoners’ Dilemma. Although it can be difficult to observe mixed strategies, our findings affirm that, when given the option to choose explicit randomized strategies, subjects are frequently willing to do so.

³⁴As we observed for initial switching rules, the mean levels are lower than the theoretical point predictions for D10 and D30, but there is no significant difference for the D0 treatment.




Figure 12. Mean switching rule for players on known incompatible ($\mu_t = 0$) platforms, prior to any observed success (Deterministic only).

Notes: Error bars indicate standard error of the mean. Dashed (red) lines indicate equilibrium point predictions.

Case (ii): Post-transgression behavior. In the second case, we distinguish between the transgressor and the compliant player after a transgression has occurred. A natural question is whether the players use the asymmetric history as a basis for subsequent asymmetric strategies after the transgression. For instance, a compliant player may hold the transgressor responsible for the resulting incompatibility and opt to remain on their current platform, while expecting the transgressor to ‘fix’ their transgression by switching back.

Figure 13a shows that the mean switching rule is indeed significantly lower for compliant players than for transgressors in all three D treatments (D0, D10, D30). These differences are reinforced by the empirical CDFs shown in the bottom row of Figure 14. There is a stark pattern of first-order stochastic dominance between the distributions of switching rules for compliant players and transgressors when they are on known incompatible platforms. For instance, in the D30 treatment, an overwhelming majority (82.6%) of compliant players chose $s = 0$ *every single time* they faced incompatibility in the subsequent periods.

A similar pattern emerges for the Stochastic treatments, although the differences between switching rules for compliant players and transgressors, shown in Figure 13b, are less stark than in the Deterministic treatments. In particular, the difference between mean switching rules for transgressors and compliant players is small in the S30 treatment. These differences are also reflected in the empirical CDFs for compliant players and transgressors. In all three Stochastic treatments, mean switching rules are very bimodal for the compliant players—a large fraction always choose $s = 0$, and a smaller, though non-trivial fraction always choose $s = 100$. Hardly any compliant players choose interior mixed strategies. This suggests that some compliant players may be trying to play one of the plausible asymmetric responses (relying on the transgressor to fix things), while others try to play the alternative plausible asymmetric response (seeking to restore compatibility themselves). Both are indicative of an attempt to use the asymmetric history as a precedent for subsequent asymmetric play. In contrast, for the transgressors, the draw towards asymmetric play is not as salient, which would explain the higher amount of mixing observed in their empirical CDFs. Furthermore, the notable difference in S30 (which exhibits a pattern of second- but not first-order stochastic dominance) is that far




Figure 13. Mean switching rule for players on known incompatible ($\mu_t = 0$) platforms following a previous transgression, separated into transgressors and compliant players.

Notes: Error bars indicate standard error of the mean. Dashed (red) lines indicate equilibrium point predictions.




Figure 14. Empirical CDFs of the switching rule for players on known incompatible ($\mu_t = 0$) platforms, after a previous transgression has occurred.

fewer transgressors always choose $s = 100$.

To summarize, in the off-equilibrium setting where the players find themselves on known incompatible platforms due to a previous transgression, we find some evidence to suggest that subjects try to use the transgression as a precedent to help them coordinate on subsequent asymmetric play.

5 Results—Coordination and Payoffs

In this section, we examine two additional outcomes of interest—coordination rates and payoffs.

Table 3. Successful coordination by treatment and by initial compatibility.

	Treatment							
	Deterministic				Stochastic			
	D0	D10	D30	Total	S0	S10	S30	Total
Overall (%)	100	100	99.07	99.62	92.86	95.54	89.42	92.65
(# pairs)	(56)	(96)	(107)	(259)	(52)	(107)	(93)	(252)
By initial platform compatibility								
<i>Incompatible (%)</i>	100	100	98.39		92.68	94.67	88.24	
(# pairs)	(40)	(60)	(61)		(38)	(71)	(60)	
<i>Compatible (%)</i>	100	100	100		93.33	97.30	91.67	
(# pairs)	(16)	(36)	(46)		(14)	(36)	(33)	

Coordination. As a first step, we document the rates of *successful coordination* (defined below) in each treatment. We say that a pair achieves *successful coordination* if they observe a success in at least one period during their match. In the Deterministic treatments, this coincides with the pair ever being on compatible platforms, while in the Stochastic treatments, it requires both platform compatibility and stochastic success in the same period. In either case, these criteria are very weak so we should expect extremely high rates of successful coordination, although it is natural to conjecture that the rates may be lower in the Stochastic treatments.

Table 3 summarizes the rates of successful coordination by treatment, and within treatment, separated by whether the pair started on compatible platforms in the first period. Consistent with the intuition given above, we observe nearly perfect rates of successful coordination in the Deterministic treatments, with only one pair failing to coordinate on compatible platforms. Furthermore, the rate of successful coordination is lower, although still very high, in the Stochastic treatments, ranging between 89.42% (in S30) and 95.54% (in S10). The bottom panel of Table 3 also shows that, as might be expected, the rate of successful coordination in Stochastic treatments is slightly lower for pairs with incompatible initial platforms.

We also examine coordination across time. In Figure 15, we plot the cumulative percentage of pairs who have observed a success over time, with each panel comparing the Deterministic and Stochastic treatments for a given switching cost. For comparison, we also include the predicted cumulative percentages (as red dashed lines) using the theoretical predictions calibrated to the parameters of the environment. The observed patterns are consistent in each panel. The fraction of pairs who have observed a success in the early periods is higher in the Deterministic treatments than in the Stochastic treatments, regardless of the cost of switching.

Payoffs. How do the differences between treatments in terms of behavior and outcomes translate to differences in payoffs? Figure 16a illustrates the mean payoffs (per period) for each treatment, using Match 1 data only, with 95% confidence intervals obtained from a cluster wild bootstrap (using each pair as an independent cluster). Mean payoffs are higher in S0 than in S10 or S30, although the difference is only statistically significant with respect to S10 ($p = 0.042$ for S0 vs. S10, and $p = 0.232$ for S0 vs. S30). On the other hand, in the Deterministic treatments, the mean payoffs are lower in D10 than in D0 or D30, although the differences fall short of statistical significance ($p = 0.103$ for D10 vs. D0, and $p = 0.102$ for D10 vs. D30).

In the other panel, Figure 16b shows that when we use data from Match 4 only, the payoff




Figure 15. Cumulative percentage of pairs who have observed a success over time.




Figure 16. Mean payoffs (per period) by treatment using (a) Match 1 only and (b) Match 4 only.

Notes: Error bars represent 95% cluster wild bootstrap confidence intervals (clustered by group).

differences for the Stochastic treatments are reversed. Mean payoffs per period are higher in S10 and S30 than in S0, although neither of the differences are statistically significant. The comparisons of mean payoffs for the Deterministic treatments are essentially the same as for Match 1. More noticeable, mean payoffs are significantly higher in Match 4 than in Match 1 for the S10 and S30 treatments. This coincides with the observation that, after gaining experience, subjects are far less likely to choose positive switching rules when they are on known compatible platforms, which leads to fewer initial transgressions in Match 4 than in Match 1. Specifically, as discussed in Section 4.2, the percentage of subjects who always choose $s = 0$ when $\mu = 1$ (prior to any transgression) increases from 50.0% to 84.5% in S10 and from 67.3% to 96.7% in S30. Fewer transgressions imply both cost savings (immediately and in future periods) and higher expected payoffs (due to maintained compatibility). Figure 16 suggests that these translate to a 16% increase in mean payoffs per period for S10 and a 13% increase in mean payoffs per period for S30.

6 Conclusion

In this paper, we study dynamic coordination between two parties who must learn *how to coordinate* over time. Unlike previous related studies, we capture various sources of friction associated with noisy outcomes and the cost of switching between platforms or procedures. We

develop and analyze a model of dynamic coordination with switching costs in which players are initially uncertain about whether or not they are on compatible platforms.

In our Deterministic setting, we derive the efficient symmetric equilibrium and provide the closed-form solution for the players' equilibrium switching probability as a function of their common belief about compatibility. Though the complexity of the dynamic problem impedes the derivation of a closed-form expression in the Stochastic setting, we devise a procedure to obtain a symmetric MPE approximation. In each setting, the optimal strategy takes a simple form, prescribing REMAIN with certainty at all beliefs above a cutoff belief (that depends on the switching cost and other parameters) and probabilistic switching otherwise. Intuitively, our model predicts that players switch with lower probability when (i) they are more optimistic about compatibility, (ii) switching is more costly, or (iii) success on compatible platforms is less likely.

To test our theoretical predictions and provide empirical validation for the model, we conducted a virtual laboratory experiment in which we varied both the payoff setting and the switching cost. Overall, the experimental findings are consistent with many of the comparative statics predictions of the model, although there are also significant departures from the theoretical predictions. Once they reach known compatible platforms, subjects choose mean switching rules close to, but statistically different from, zero. As a result, we observe a pattern of transgressions that shift pairs away from (known) compatible platforms. We argue that these disparities between theory and data may be explained by appealing to cognitive limitations, or to a form of recency bias, whereby inattention causes players to focus on the number of consecutive failures, rather than the underlying state. In a similar fashion, subjects' sensitivity to the number of consecutive failures can help to rationalize the attenuated responses to intermediate beliefs observed in our costly Stochastic treatments.

After a transgression, those who are ostensibly responsible (the transgressors) tend to choose switching rules that are substantially higher than their compliant counterparts. Thus, while transgressions can be destructive, players tend to use the identity of the transgressor as a basis for subsequent attempts to re-establish coordination using asymmetric play. Successful coordination eventually occurs with high frequency in all treatments, but at a slightly lower rate in the Stochastic setting, as predicted by the theory. Still, whereas average payoffs (per period) are initially higher in the simpler, Deterministic treatments than in the Stochastic treatments, the differences in payoffs between settings disappear as subjects accrue experience dealing with the costly, noisy environment.

Naturally, there are several directions in which our research may be extended. The theoretical model that we present is relatively stylized, focusing on a strategic interaction between two players who each have two possible platforms. Coordination within and across organizations often relies on the efforts of many agents who interact with varied frequency and have potentially many possible configurations of procedures from which to choose. In future work, it may be interesting to generalize the framework to allow for more players, larger state and action spaces, and complex interaction structures (e.g., using a network model) to explore more of the intricacies present in real-world organizations. An additional, related extension might involve allowing for different levels of success based on the players' degrees of compatibility. However, such generalizations may come at the cost of tractability.

By implementing the above theoretical extensions and introducing other interventions, future experiments could provide further insight into coordination in decentralized organizations.

In our experiment, we observe some “under-switching” (relative to the point predictions) at low beliefs and some anomalous switching in pairs that have successfully coordinated. In future work, it may be important to consider the extent to which under-switching is driven by non-Bayesian updating of the common belief when failure is not perfectly informative. Furthermore, many previous experimental studies have shown that communication can often improve coordination. While it is unclear that introducing communication into this framework would improve pairs’ ability to coordinate in the first place (especially in the more complicated Stochastic setting), communication could provide more sophisticated players with a tool to teach less sophisticated players how to switch optimally and how to *stay* coordinated. Future work can address whether players, in practice, use communication effectively to overcome the strategic tensions of our environment.

Declaration of competing interest The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability All data and code required to replicate the findings of this study are available in the Open Science Framework (OSF). The repository contains the raw data, processed data, Stata .do files, and z-Tree programs, accessible via <https://doi.org/10.17605/OSF.IO/NP84C>.

ONLINE APPENDIX

Dynamic coordination with switching costs

Joe S. Ballard Luke Boosey

A General Theoretical Model

A.1 Preliminaries

Consider an infinite-horizon coordination game with two players, $i \in \{1, 2\}$, in discrete time, indexed by $t = 0, 1, 2, \dots$. At the beginning of time, Nature randomly and privately assigns each player i an initial platform $\theta_{i,0} \in \Theta_i = \{\theta_i^1, \theta_i^2\}$. While it is common knowledge that there are two possible platforms in each of Θ_1 and Θ_2 , neither player can distinguish between the other player's platforms. Furthermore, each player's current platform remains her private information for the entire game, as we assume that communication between players is impossible.

Nature randomly draws two (out of four) platform pairings to be *compatible* for the players. For players i, j and platforms k, l , if (θ_i^k, θ_j^l) is a compatible pairing, then (θ_i^l, θ_j^k) is also compatible, i.e., each of i 's platforms are compatible with exactly one of j 's platforms. Compatible platform pairings are fixed for the entire interaction, though neither player is ever informed about which platform pairings are compatible. Coordinating on a compatible platform pairing results in a higher (expected) benefit for both players, which is described in greater detail below.

We assume that compatible pairings do not differ in their productivity; thus, it is convenient to denote the period- t state of the world by $\omega_t \in \{C, I\}$, where players are either on a compatible (C) platform pairing or an incompatible (I) one. Based on the draw of Nature, the initial state of the world is ω_0 , which is not observed by either player. At the beginning of the game, players hold a common prior $\mu_0 \in (0, 1)$ that they are on a compatible pairing, $\omega_0 = C$.

In each period t , the state of the world transitions based on the actions taken by the players. Players then receive a benefit that serves as a (potentially noisy) public signal about compatibility, and consequently, update their (common) belief about the state. Depending on the setting and the realized outcome, players may be able to infer the current state of the world; however, neither player ever *explicitly* observes whether they are on a compatible platform pairing.³⁵

A.2 Timing

The timing of each period t in the dynamic coordination game is as follows: players enter period t with common prior μ_t about the unobservable state ω_t . Then, they move simultaneously, with each player i choosing an action $a_{i,t} \in A = \{R, S\}$ to either REMAIN (R) on their current platform or SWITCH (S) to the other platform. Choosing S has associated fixed cost $c(S) = c \geq 0$, while choosing R is costless, $c(R) = 0$. We represent the players' period- t action profile by $a_t = (a_{1,t}, a_{2,t})$, which is observed by both players after their actions are selected.

Based on the players' actions, the state ω_t transitions to $\hat{\omega}_t$. In particular, state $\hat{\omega}_t =$

³⁵In particular, players will be able to infer the current state of the world perfectly in the Deterministic setting (i.e., following either success or failure) and imperfectly in the Stochastic setting (i.e., only following a success).

$\psi(a_t, \omega_t)$ is determined via the following transition rule:

$$\psi(a, \omega) = \begin{cases} C, & \text{if } (a, \omega) \in \left\{ \left((S, S), C \right), \left((R, R), C \right), \left((S, R), I \right), \left((R, S), I \right) \right\}, \\ I, & \text{otherwise.} \end{cases}$$

Thus, the state in our model is endogenous, but it evolves deterministically over time. If the players begin the period on a compatible (incompatible) platform pairing, they remain on a compatible (incompatible) pairing if they choose the same action, i.e., both REMAIN or both SWITCH, but move to an incompatible (compatible) pairing if they choose different actions, i.e., one REMAIN and one SWITCH.

Following the state transition from ω_t to $\hat{\omega}_t$, players each receive a common benefit $y_t \in \{y^L, y^H\}$ that depends on (1) whether their resulting platforms are compatible and (2) the draw of a random variable. If platforms are incompatible, each player receives benefit $y^L \geq 0$ (“failure”) with certainty; if platforms are compatible, each player receives benefit $y^H > y^L$ (“success”) with probability $p \in (0, 1]$ and benefit y^L with probability $1 - p$. The period- t payoff for player i is then $\pi_{i,t}(a_t, y_t) = y_t - c(a_{i,t})$, where $c(a_{i,t})$ is player i ’s individual cost as a function of her chosen action.

After observing actions and receiving payoffs for period t , players update their beliefs about compatibility. In particular, we assume that players update their beliefs according to Bayes’ rule. Since players have the same prior and observe the same information in each period, this process leads to identical posterior beliefs $\mu_{1,t+1} = \mu_{2,t+1} = \mu_{t+1}$.³⁶ Because state transitions are deterministic, given μ_t , a_t , and y_t , the common posterior belief μ_{t+1} is

$$\mu_{t+1} = \begin{cases} \frac{\mu_t(1-p)}{1-p\mu_t}, & \text{if } (a_t, y_t) \in \left\{ \left((S, S), y^L \right), \left((R, R), y^L \right) \right\}, \\ \frac{(1-\mu_t)(1-p)}{1-p(1-\mu_t)}, & \text{if } (a_t, y_t) \in \left\{ \left((S, R), y^L \right), \left((R, S), y^L \right) \right\}, \\ 1, & \text{if } y_t = y^H. \end{cases}$$

Because players’ actions directly impact how μ_t updates, for clarity, we will often use notation that indicates the updated belief conditional on whether the players choose the same or different actions. That is, let μ^+ represent μ_{t+1} following $(a_i = a_j, y^L)$, and let μ^- represent μ_{t+1} following $(a_i \neq a_j, y^L)$. After belief updating, play proceeds to period $t + 1$, where the state $\omega_{t+1} = \hat{\omega}_t$ carries over based on the actions of period t and players hold common belief μ_{t+1} . Figure A.1 summarizes the timing of each period t .




Figure A.1. Period- t timing

³⁶Due to this construction, our model can be viewed alternatively as a stochastic game with observable states, where the players’ common belief is the state variable (see, e.g., ?).

A.3 Solution Concept

Our solution concept is Markov perfect equilibrium (MPE), a refinement of subgame perfect equilibrium that is often applied to the study of stochastic games. Within the set of MPE, we focus on symmetric MPE in which players choose identical strategies that depend only on their commonly held belief μ .³⁷ A mixed Markov strategy for player i is a function s_i that assigns to each $\mu \in [0, 1]$ the probability with which player i chooses SWITCH. A profile of Markov strategies that constitutes a subgame perfect equilibrium of the dynamic coordination game is called a Markov perfect equilibrium, and a symmetric MPE is one in which $s_i(\mu) = s_j(\mu)$ for all μ .

As in many other settings that model coordination games and dynamic interactions, restricting attention to symmetric strategies still results in a multiplicity of equilibria in our environment. To obtain a sharper prediction for our laboratory experiment, we argue that the most salient symmetric MPE is one in which both players choose REMAIN with certainty for all $\mu \geq \mu^*$ and choose SWITCH with some probability $s(\mu) \in (0, 1)$ for all $\mu < \mu^*$, where μ^* is an equilibrium cutoff belief that depends on the switching cost and other parameters of the model. Additionally, for $c > 0$, this equilibrium has the advantage of being the *efficient* symmetric MPE, as it both maximizes the players' probability of achieving successful coordination in each period and minimizes the costs incurred as a result of switching platforms.

Throughout the analysis, it is convenient to use a value function $V : [0, 1] \rightarrow \mathbb{R}$ to characterize our equilibrium. Specifically, $V(\mu)$ denotes each player's expected discounted payoff stream ("continuation value") at the beginning of a period in which the common prior equals μ . Extending our above notation for the updated belief conditional on whether players choose the same or different actions, the corresponding continuation values are defined as

$$V^+ := V(\mu_{t+1} | a_i = a_j, y^L) \quad \text{and} \quad V^- := V(\mu_{t+1} | a_i \neq a_j, y^L).$$

To guarantee the existence of equilibria in which players choose SWITCH with positive probability at *some* belief, we maintain the following parameter assumption:

Assumption 1. Parameters δ , c , p , y^H , and y^L are such that

$$c < \frac{p(y^H - y^L)}{1 - \delta}. \tag{A.1}$$

Note that $\frac{y^L}{1 - \delta}$ is the minimum of the set of continuation values, as each player can guarantee herself at least y^L in every period by simply choosing REMAIN with certainty. Furthermore, $\frac{py^H + (1 - p)y^L}{1 - \delta}$ is the maximum of the set of continuation values because the expected period-payoff from being on a compatible platform pairing is $py^H + (1 - p)y^L$. Assumption 1 ensures that the cost of switching is strictly lower than the difference between these extreme values, i.e., the largest possible increase in the expected payoff stream that results from achieving coordination.

A.4 Deterministic Setting, $p = 1$

We begin our analysis with an examination of the Deterministic setting. In this setting, because players are guaranteed success when they are on a compatible platform pairing, $p = 1$, the state of the world ω is fully revealed by the realized benefit y . Here, the game takes a rather simple

³⁷See Section 2.3 for a discussion of our focus on symmetric Markov strategies.

form, where players choose actions in period 0 based on μ_0 and update their common belief to either $\mu_1 = 0$ or $\mu_1 = 1$, conditional on y_0 . In every subsequent period, players are able to track the underlying state due to the perfect observability of actions and state-revealing benefits. As a result, there are only three unique values of μ that can occur in equilibrium: $\mu_0 \in (0, 1)$, $\mu = 0$, and $\mu = 1$.

Consider the case in which players observe a success in period $t - 1$. Following $y_{t-1} = y^H$, players update their belief via Bayes' rule to $\mu_t = 1$. At this belief, the equilibrium prediction involves both players choosing REMAIN forever. To see why, observe that, conditional on player j choosing REMAIN in every period t where $\mu_t = 1$, it is also optimal for player i to choose REMAIN in period t . If players are currently on a compatible platform pairing, continued play of (REMAIN, REMAIN) maximizes joint benefits in the continuation game as it ensures that players stay on a compatible platform pairing for the rest of the game and, consequently, receive y^H in every future period. At $\mu_t = 1$, players can also maintain compatibility with (SWITCH, SWITCH) in every period; however, if $c > 0$, this equilibrium behavior results in repeated payment of the switching cost by both players and is thus inefficient. Even in the absence of switching costs, $c = 0$, (SWITCH, SWITCH) is arguably less salient than remaining on the established compatible platform pairing. Thus, in the Deterministic setting, when $\mu = 1$, the efficient symmetric MPE prescribes play of REMAIN with certainty for each player if $c > 0$. While it is not possible to select this equilibrium over one involving (SWITCH, SWITCH) on the basis of efficiency when $c = 0$, we argue that (REMAIN, REMAIN) is the most natural prescription of equilibrium play as it preserves the status quo. In equilibrium, players are able to maintain compatibility once success is realized, resulting in a corresponding continuation value of $V(1) = \frac{y^H}{1-\delta}$.

Consider now the case in which players enter period t with belief $\mu_t < 1$. This case encompasses period 0, where players begin the game with common prior $\mu_0 \in (0, 1)$, as well as each period $t > 0$ that follows a failure in period $t - 1$, $y_{t-1} = y^L$, where players know that they are currently on incompatible platforms, $\mu_t = 0$. Given belief μ_t , continuation values $V(\mu_{t+1})$, and player j 's switching probability $s_{j,t}$, player i 's expected payoff from choosing SWITCH in period t is

$$(\mu_t s_{j,t} + (1 - \mu_t)(1 - s_{j,t})) \left[y^H + \delta \left(\frac{y^H}{1 - \delta} \right) \right] + (\mu_t(1 - s_{j,t}) + (1 - \mu_t)s_{j,t}) \left[y^L + \delta V(0) \right] - c,$$

while player i 's expected payoff from choosing REMAIN is

$$(\mu_t(1 - s_{j,t}) + (1 - \mu_t)s_{j,t}) \left[y^H + \delta \left(\frac{y^H}{1 - \delta} \right) \right] + (\mu_t s_{j,t} + (1 - \mu_t)(1 - s_{j,t})) \left[y^L + \delta V(0) \right].$$

Thus, it is optimal for player i to choose SWITCH if and only if

$$(1 - 2s_{j,t})(1 - 2\mu_t) \left[y^H - y^L + \delta \left(\frac{y^H}{1 - \delta} - V(0) \right) \right] \geq c, \quad [\text{A.2}]$$

with indifference between SWITCH and REMAIN at equality.

From (A.2), one observation is immediate. Clearly, the expression in square brackets is strictly positive because $\frac{y^H}{1-\delta}$ is the maximum of the set of continuation values and $y^H > y^L$. This implies that, for sufficiently low values of μ_t , the only symmetric profile that satisfies (A.2) must involve both players choosing SWITCH with some probability $s(\mu_t) \in (0, 1)$. In the efficient symmetric MPE, there exists an equilibrium cutoff belief $\mu^{D*} \in (0, \frac{1}{2}]$ such that players choose

SWITCH with probability $s(\mu_t) \in (0, 1)$ for all $\mu_t < \mu^{D*}$; however, at this point in our analysis, it is not obvious whether μ^{D*} exists, and if so, what its value should be. For the time being, we conjecture the existence of such a cutoff belief and derive optimal behavior.

Suppose that $\mu_t < \mu^{D*}$. By imposing symmetry and solving a binding (A.2), we obtain $s(\mu_t)$ as a function of the continuation value $V(0)$:

$$s(\mu_t) = \frac{1}{2} \left\{ 1 - \frac{c}{(1 - 2\mu_t) \left(\frac{y^H}{1-\delta} - y^L - \delta V(0) \right)} \right\}. \quad [\text{A.3}]$$

Once their common belief reaches $\mu = 0$, the players face an identical problem in each period, conditional on continued failure. In particular, we can write $V(0)$ in the following stationary form, as a function of equilibrium strategies, where $s(0)$ equals (A.3) evaluated at $\mu = 0$:

$$V(0) = \left(1 - 2s(0)(1 - s(0)) \right) \left(y^L + \delta V(0) \right) + 2s(0)(1 - s(0)) \left(\frac{y^H}{1-\delta} \right) - s(0)c.$$

That is, in the efficient symmetric MPE, when players know that they are on an incompatible platform pairing in a given period, they remain uncoordinated for that period (i.e., stay on an incompatible pairing) with probability $1 - 2s(0)(1 - s(0))$, but obtain successful coordination (i.e., switch to a compatible pairing) with probability $2s(0)(1 - s(0))$. After some simplification, our final expression for $V(0)$ is

$$V(0) = \frac{y^H + (1 - \delta)(y^L - c)}{(2 - \delta)(1 - \delta)}, \quad [\text{A.4}]$$

which, by Assumption 1, is strictly greater than the lower bound of the set of continuation values, $\frac{y^L}{1-\delta}$. Equilibrium switching probabilities $s(\mu_t)$ for $\mu_t < \mu^{D*}$ are obtained by substituting (A.4) back into (A.3):

$$s(\mu_t) = \frac{1}{2} \left[1 - \frac{(1 - \frac{\delta}{2})c}{(1 - 2\mu_t)(y^H - y^L + \frac{\delta}{2}c)} \right]. \quad [\text{A.5}]$$

A number of interesting features are worth noting about $s(\mu_t)$. If $c = 0$, then $s(\mu_t) = \frac{1}{2}$ for all $\mu_t < \mu^{D*}$. In this case, the optimal switching probability is that which maximizes the players' probability of achieving successful coordination in the current period, $\mu_t(1 - 2s_t(1 - s_t)) + (1 - \mu_t)(2s_t(1 - s_t))$. Since $\mu_t < \mu^{D*} \leq \frac{1}{2}$, this probability is strictly increasing in $2s_t(1 - s_t)$, a term which is maximized by $s_t = \frac{1}{2}$.

If $c > 0$, players optimally choose SWITCH with probability strictly less than $\frac{1}{2}$. Notably, $s(\mu_t)$ depends on μ_t , c , $y^H - y^L$, and δ . In this case, clear comparative statics emerge. For a given $\mu_t < \mu^{D*}$, $s(\mu_t)$ is strictly increasing in $y^H - y^L$, strictly increasing in δ , and strictly decreasing in c . If instead we allow μ_t to vary within $[0, \mu^{D*})$ and hold all other parameters fixed, then $s(\mu_t)$ is strictly decreasing in μ_t , as well. Intuitively, players are more willing to switch when the benefits from achieving coordination – both in the short-term ($y^H - y^L$) and long-term (δ) – are greater. Similarly, they are less willing to switch when it is more costly to do so. As players become more certain that they are on incompatible platforms, they switch with higher probability, but the equilibrium rate remains below $\frac{1}{2}$ due to the presence of switching costs.

To derive the equilibrium cutoff belief μ^{D*} , we set $s(\mu^{D*})$ equal to zero and solve for μ^{D*} ,

which gives us

$$\mu^{D*} = \frac{1}{2} \left[1 - \frac{(1 - \frac{\delta}{2})c}{y^H - y^L + \frac{\delta}{2}c} \right].$$

Since $y^H > y^L$ and $c \geq 0$, $\mu^{D*} \leq \frac{1}{2}$. Furthermore, by Assumption 1, $\mu^{D*} > 0$. As conjectured, there exists an equilibrium cutoff belief $\mu^{D*} \in (0, \frac{1}{2}]$ such that players choose SWITCH with probability $s(\mu_t) \in (0, 1)$ for all $\mu_t < \mu^{D*}$.

Equipped with our equilibrium cutoff belief μ^{D*} , we discuss optimal behavior for all beliefs $\mu_t \in [\mu^{D*}, 1)$. Substituting the derived value of $V(0)$ into (A.2), we observe that player i prefers to choose SWITCH if and only if

$$(1 - 2s_{j,t})(1 - 2\mu_t) \left[\frac{y^H - y^L + \frac{\delta}{2}c}{1 - \frac{\delta}{2}} \right] \geq c, \quad [\text{A.6}]$$

with indifference between SWITCH and REMAIN at equality. As previously noted, the expression in square brackets is strictly positive. In fact, by Assumption 1, $\left[\frac{y^H - y^L + \frac{\delta}{2}c}{1 - \frac{\delta}{2}} \right] > c$, which is RHS(A.6). However, for $\mu_t \geq \mu^{D*}$, $(1 - 2s_{j,t})(1 - 2\mu_t)$ may be positive, negative, or equal to zero, depending on $s_{j,t}$ and μ_t . As a result, equilibrium selection requires some care and is most easily undertaken by analyzing the cases of $c > 0$ and $c = 0$ separately.

When switching is costly, $c > 0$, the equilibrium cutoff belief is $\mu^{D*} < \frac{1}{2}$. For $\mu_t \in [\mu^{D*}, \frac{1}{2})$, (REMAIN, REMAIN) is the *only* symmetric equilibrium profile, and for $\mu_t = \frac{1}{2}$, REMAIN is the dominant action for both players. Both of these facts can be seen from the simplified optimality condition, where, given $s_{j,t} = 0$, LHS(A.6) decreases from c to 0 as μ_t increases from μ^{D*} to $\frac{1}{2}$. For $\mu_t \in (\frac{1}{2}, 1)$, it is possible to support play of (REMAIN, REMAIN), (SWITCH, SWITCH), or even mixing between REMAIN and SWITCH with identical probabilities in a symmetric equilibrium. As noted earlier, the players' probability of success in each period is $\mu_t(1 - 2s_t(1 - s_t)) + (1 - \mu_t)(2s_t(1 - s_t))$, which is decreasing in $2s_t(1 - s_t)$ since $\mu_t > \frac{1}{2}$. It is thus maximized by either $s_t = 0$ (REMAIN with certainty) or $s_t = 1$ (SWITCH with certainty), but efficiency requires that both players choose REMAIN with certainty for $\mu_t \in (\frac{1}{2}, 1)$ as this minimizes their joint costs from switching. Thus, the efficient symmetric MPE prescribes play of REMAIN with certainty for both players for all $\mu_t \geq \mu^{D*}$.

When switching is costless, $c = 0$, the equilibrium cutoff belief is $\mu^{D*} = \frac{1}{2}$, and there exists a symmetric MPE in which both players choose REMAIN with certainty for all $\mu_t \geq \mu^{D*}$; however, there exist additional symmetric MPE in which both players choose SWITCH with certainty or both mix between REMAIN and SWITCH with probability $\frac{1}{2}$ over the same belief range. Unlike the case where $c > 0$, in the absence of switching costs, multiple symmetric equilibria maximize the players' probability of success in each period, and no inefficiencies can result from joint switching while players are on compatible platforms. Consequently, there are multiple efficient symmetric MPE and thus no unique prediction for equilibrium behavior. One could argue that players who are more optimistic about compatibility than not ($\mu_t > \frac{1}{2}$) should remain where they are, due to the salience of the status quo alternative. On the other hand, given that a success has not yet occurred, one could argue that the players' current platform pairing is not yet established as a focal point. To maintain consistency with our analysis in the costly switching case, we hold the former position.

We summarize our equilibrium for the Deterministic setting in Proposition 1:

Proposition 1. *Consider the Deterministic setting of the dynamic coordination game, where $p = 1$, and suppose that Assumption 1 holds. Then there exists a symmetric MPE in which, given common belief μ_t , the equilibrium probability of switching in period t is*

$$s^{D*}(\mu_t) = \begin{cases} \frac{1}{2} \left[1 - \frac{(1 - \frac{\delta}{2})c}{(1 - 2\mu_t)(y^H - y^L + \frac{\delta}{2}c)} \right], & \text{if } \mu_t < \mu^{D*}, \\ 0, & \text{otherwise,} \end{cases}$$

with the corresponding equilibrium cutoff belief,

$$\mu^{D*} = \frac{1}{2} \left[1 - \frac{(1 - \frac{\delta}{2})c}{y^H - y^L + \frac{\delta}{2}c} \right].$$

If $c > 0$, then this equilibrium is the efficient symmetric MPE.

The equilibrium described in Proposition 1 provides us with a straightforward prediction for behavior in the Deterministic setting. In period 0, switching behavior is determined by the prior μ_0 . If $\mu_0 \geq \mu^{D*}$, both players choose REMAIN with certainty. If $\mu_0 < \mu^{D*}$, players mix between SWITCH and REMAIN according to $s^{D*}(\mu_0)$, placing higher probability on SWITCH as μ_0 decreases towards zero if $c > 0$. In the absence of switching costs, $c = 0$, players mix between REMAIN and SWITCH with equal probability ($\frac{1}{2}$) for all $\mu_t < \mu^{D*}$. In every period $t > 0$, if a success was obtained in period $t - 1$, both players choose REMAIN with certainty (and thus maintain coordination forever). Otherwise, knowing that failure indicates their incompatibility, players mix between SWITCH and REMAIN according to $s^{D*}(0)$.

Switching costs influence optimal behavior through two main channels, provided that $c < \frac{p(y^H - y^L)}{1 - \delta}$. First, given $\mu_t < \mu^{D*}$, a higher cost of switching induces a strictly lower probability of switching. Second, increasing c leads to a strictly lower cutoff belief μ^{D*} . Overall, when it becomes more costly for players to switch platforms, switching occurs less often, over a smaller range of only the most pessimistic beliefs.

A.5 Stochastic Setting, $p \in (0, 1)$

We now turn our attention to the Stochastic setting. As in the Deterministic setting, players are able to infer the state ω once a success y^H has occurred. From that point forward, players can track the underlying state due to the perfect observability of actions, resulting in a common belief of either $\mu = 1$ or $\mu = 0$ in each period. However, since $p \in (0, 1)$, success is not guaranteed when players are on a compatible platform pairing. As a result, until success occurs, players faced with repeated failure are unable to distinguish incompatibility from a string of “bad luck.” This limited ability to infer the state based on the resulting benefit compounds the existing tension due to switching costs, as the players’ decisions involve coordinating not only *who should switch* but also *whether anyone should switch*.

The strategic environment is thus necessarily more complex owing to the evolution of the common belief prior to successful coordination. If players have yet to achieve success and choose the same action, both REMAIN or both SWITCH, then failure causes their belief to deteriorate towards zero, $\mu^+ < \mu_t$. However, players that face failure and choose different actions, one REMAIN and one SWITCH, may become more or less optimistic about being on compatible platforms, depending on their current belief. If $\mu_t \geq \frac{1}{2}$, then observing $a_{i,t} \neq a_{j,t}$ leads players

to doubt compatibility, and their belief erodes even further following y^L , $\mu^- < \mu_t$. If instead $\mu_t < \frac{1}{2}$, players believe compatibility is more likely than not after $a_{i,t} \neq a_{j,t}$, but this optimism is dampened by y^L , $\mu^- > \mu_t$.

Starting from a generic prior $\mu_0 \in (0, 1)$, the players remain uncertain about the state of the world until a success occurs. For example, players become more pessimistic about compatibility if they remain on the same platforms and repeatedly fail, but their belief never reaches zero. Furthermore, unless their prior takes a particular form (which we define in Appendix B), it is unlikely that they hold the same interior belief in multiple periods of the game, even if they switch to likely compatible platforms and remain on those platforms through a sequence of failures. That is, in general, μ does not cycle through values within a fixed set. For these reasons, there are infinitely many beliefs that can arise in equilibrium.

Consider first the case in which players' common belief at the beginning of period t is $\mu_t = 1$, i.e., the case in which players observe a success in period $t - 1$, $y_{t-1} = y^H$. For reasons identical to those discussed in Appendix A.4, our equilibrium prediction involves both players choosing REMAIN forever at this belief. This prescription of play maximizes joint expected benefits while also minimizing joint costs from switching, with strict efficiency gains relative to other symmetric equilibria that prescribe (SWITCH, SWITCH) or mixing between SWITCH and REMAIN if $c > 0$. On compatible platforms, players' expected period-benefit is $py^H + (1 - p)y^L$, so the corresponding continuation value is $V(1) = \frac{py^H + (1 - p)y^L}{1 - \delta}$.

Consider instead the case in which players enter period t with common belief $\mu_t < 1$, which encompasses period 0, where $\mu_0 \in (0, 1)$, as well as each period that follows a failure in period $t - 1$. Given common belief μ_t , continuation values $V(\mu_{t+1})$, and player j 's switching probability $s_{j,t}$, player i 's expected payoff from choosing SWITCH in period t is

$$\begin{aligned} p(\mu_t s_{j,t} + (1 - \mu_t)(1 - s_{j,t})) \left(y^H + \delta V(1) \right) &+ (1 - s_{j,t})(1 - p(1 - \mu_t)) \left(y^L + \delta V^- \right) \\ &+ s_{j,t}(1 - p\mu_t) \left(y^L + \delta V^+ \right) - c, \end{aligned}$$

while player i 's expected payoff from choosing REMAIN is

$$\begin{aligned} p(\mu_t(1 - s_{j,t}) + (1 - \mu_t)s_{j,t}) \left(y^H + \delta V(1) \right) &+ s_{j,t}(1 - p(1 - \mu_t)) \left(y^L + \delta V^- \right) \\ &+ (1 - s_{j,t})(1 - p\mu_t) \left(y^L + \delta V^+ \right). \end{aligned}$$

Thus, it is optimal for player i to choose SWITCH in period t if and only if

$$(1 - 2s_{j,t}) \left\{ p(1 - 2\mu_t)(y^H - y^L) + \delta \left[p(1 - 2\mu_t)V(1) + (1 - p(1 - \mu_t))V^- - (1 - p\mu_t)V^+ \right] \right\} \geq c,$$

with indifference between SWITCH and REMAIN at equality. For the arguments that follow, it is more convenient to reference a compact version of this condition,

$$(1 - 2s_{j,t})\Lambda(\mu_t) \geq c, \tag{A.7}$$

where $\Lambda(\mu_t) = p(1 - 2\mu_t)(y^H - y^L) + \delta \left[p(1 - 2\mu_t)V(1) + (1 - p(1 - \mu_t))V^- - (1 - p\mu_t)V^+ \right]$.

This optimality condition shares similarities with the one derived in the Deterministic set-

ting, but there are also some important differences. In general, each $\mu_t \in (0, 1)$ updates in one of three ways, conditional on the players' actions and their realized benefit. If success occurs, then players update their belief to $\mu_{t+1} = 1$ as in the Deterministic setting; however, failure leads players to a posterior belief of either μ^+ (if players choose the same actions) or μ^- (if players choose different actions). The probability of each of these paths, and thus the players' optimal behavior, relies on not only the belief μ_t but also the success probability p . As $p \rightarrow 1$, the optimality condition closely resembles that of the Deterministic setting.³⁸ As $p \rightarrow 0$, the net expected benefit from achieving successful coordination goes to zero, in which case players always strictly prefer to choose REMAIN with certainty if $c > 0$ and weakly prefer to do so if $c = 0$. The more interesting case lies at intermediate values of p ; maintaining Assumption 1, we have $p > \frac{(1-\delta)c}{y^H - y^L}$.

We derive optimal behavior at three key beliefs, $\mu = 1$, $\mu = \frac{1}{2}$, and $\mu = 0$, to provide insight into the problem for a general belief $\mu_t \in (0, 1)$. As discussed above, our equilibrium predicts (REMAIN, REMAIN) at $\mu = 1$. This can also be argued from the optimality condition, as $\Lambda(1)$ is negative, indicating that a symmetric MPE can involve play of (REMAIN, REMAIN), (SWITCH, SWITCH), or both players mixing between REMAIN and SWITCH. Of these three prescriptions of equilibrium play, (REMAIN, REMAIN) is efficient if $c > 0$. For $\mu_t = \frac{1}{2}$, the probability of achieving successful coordination in the current period is $\frac{1}{2}$, and the posterior belief following failure is the same no matter which actions are taken by the players, $\mu^+ = \mu^- = \frac{1-p}{2-p}$. As a result, $\Lambda(\frac{1}{2}) = 0$, which implies that REMAIN with certainty is a dominant action for $c > 0$ and can be supported as part of a symmetric MPE for all c . Finally, for $\mu_t = 0$, the only symmetric profile that satisfies (A.7) must involve both players choosing SWITCH with some probability $s(\mu_t) \in (0, 1)$, since $\Lambda(0)$ is positive.³⁹

Notice that increasing μ_t from 0 to $\frac{1}{2}$ to 1 results in a decrease in $\Lambda(\mu_t)$ from a positive value to 0 to a negative value. This mirrors the relationship between the same three beliefs and the expression in square brackets in (A.2). Furthermore, while the difference $\delta \left[(1 - p(1 - \mu_t))V^- - (1 - p\mu_t)V^+ \right]$ may be positive, negative, or equal to zero, much of the weight in $\Lambda(\mu_t)$ can be attributed to the term $p(1 - 2\mu_t) \left[y^H - y^L + \delta V(1) \right]$, which is strictly decreasing in μ_t . Thus, even though the belief evolves in a more complicated fashion in the Stochastic setting, the form of the optimality condition suggests that players may use an equilibrium cutoff strategy similar to that derived in the Deterministic setting.

Specifically, we conjecture that there exists a symmetric MPE with an equilibrium cutoff belief $\mu^{S*} \in (0, \frac{1}{2}]$ such that both players choose SWITCH with probability $s(\mu_t) \in (0, 1)$ for all $\mu_t < \mu^{S*}$ and choose REMAIN with certainty for all $\mu_t \geq \mu^{S*}$. For $\mu_t < \mu^{S*}$, the candidate equilibrium probability of switching is found by imposing symmetry and solving a binding (A.7) for $s(\mu_t)$:

$$s(\mu_t) = \frac{1}{2} \left[1 - \frac{c}{\Lambda(\mu_t)} \right]. \quad [\text{A.8}]$$

³⁸As expected, substituting $p = 1$ into (A.7) results in an optimality condition that is *exactly the same* as the one derived in Appendix A.4.

³⁹In the Stochastic setting, $\mu = 0$ can only be reached by players who have achieved successful coordination and subsequently chosen different actions. Based on our prediction, players should never reach $\mu = 0$ on the equilibrium path. However, because Markov strategies only depend on the payoff-relevant state, μ , and not the full history of the game, past transgressions or mistakes do not play a role in an MPE. Thus, in the event that players reach $\mu = 0$, equilibrium behavior is still governed by the above optimality condition.

Given μ_t , $\Lambda(\mu_t)$ is a function of up to three continuation values, so calculating the equilibrium switching probability for each μ_t requires solving the system of continuation values $(V(\mu_t))_{\mu_t \in [0,1]}$.⁴⁰ The general form of $V(\mu_t)$ can be written as

$$\begin{aligned} V(\mu_t) = & 2s_t(1-s_t) \left\{ p(1-\mu_t) \left(y^H + \delta V(1) \right) + (1-p(1-\mu_t)) \left(y^L + \delta V^- \right) \right\} \\ & + (1-2s_t(1-s_t)) \left\{ p\mu_t \left(y^H + \delta V(1) \right) + (1-p\mu_t) \left(y^L + \delta V^+ \right) \right\} \\ & - s_t c, \end{aligned}$$

where s_t represents the equilibrium switching probability at μ_t . This expression can be simplified based on whether μ_t lies above or below the equilibrium cutoff belief μ^{S*} . That is, for $\mu_t \geq \mu^{S*}$,

$$V(\mu_t) = p\mu_t \left(y^H + \delta V(1) \right) + (1-p\mu_t) \left(y^L + \delta V^+ \right),$$

while for $\mu_t < \mu^{S*}$,

$$V(\mu_t) = \frac{1}{2} \left\{ py^H + (2-p)y^L - c + \delta \left[pV(1) + (1-p(1-\mu_t))V^- + (1-p\mu_t)V^+ \right] \right\}.$$

Using our value for $V(1)$ and the candidate equilibrium, we can immediately solve for $V(0)$:

$$\begin{aligned} V(0) &= \frac{1}{2} \left\{ py^H + (1-p)y^L + y^L - c + \delta \left[V(1) + V(0) \right] \right\} \\ \iff V(0) &= \frac{py^H + (1-p)y^L + (1-\delta)(y^L - c)}{(2-\delta)(1-\delta)}. \end{aligned}$$

Equilibrium switching rate $s(0)$ is then obtained by substituting $\mu = 0$, $V(1)$, and $V(0)$ back into (A.8). Note that $V(0) > \frac{y^L}{1-\delta}$ by Assumption 1, so mixing with probability $s(0)$ at $\mu = 0$ is indeed optimal.

Through $V(1)$ and $V(0)$, we have pinned down optimal behavior at the boundary beliefs, $\mu = 1$ and $\mu = 0$. However, because there are infinitely many beliefs that can arise in equilibrium, there are infinitely many corresponding continuation values to consider, as well. To combat the difficulties associated with solving an infinite system of simultaneous linear equations, we devise a method to compute an equilibrium approximation for the Stochastic setting. Our approximation method centers on implementing the structure of our candidate equilibrium and limiting the size of the set of continuation values. By constructing a *finite* system of continuation values, we can exploit the recursive nature of each continuation value and compute a solution of the system using standard tools of linear algebra. These continuation values can then be used in (A.8) to calculate the switching probability for each μ in our restricted set. Precise details of our method are given in Appendix B. Here, we comment on our equilibrium approximation and provide some intuition.

Consider the Stochastic setting of the dynamic coordination game, where $p \in (0, 1)$, and

⁴⁰Note that, if $c = 0$, $s(\mu_t) = \frac{1}{2}$ for all $\mu_t < \mu^{S*}$, i.e., optimal behavior at low beliefs involves mixing between REMAIN and SWITCH with equal probability. In the absence of switching costs, players simply maximize the probability of achieving successful coordination in the current period, $p \left[\mu_t (1 - 2s_t(1-s_t)) + (1-\mu_t)(2s_t(1-s_t)) \right]$. Since $\mu_t < \mu^{S*} \leq \frac{1}{2}$, this probability is maximized by $s_t = \frac{1}{2}$, as in the Deterministic setting.

suppose that Assumption 1 holds. Then there exists a symmetric MPE approximation in which, given common belief μ_t , the probability of switching in period t is

$$s^{S*}(\mu_t) = \begin{cases} \frac{1}{2} \left[1 - \frac{c}{\Lambda(\mu_t)} \right], & \text{if } \mu_t < \mu^{S*}, \\ 0, & \text{otherwise,} \end{cases}$$

where $\Lambda(\mu_t) = p(1 - 2\mu_t)(y^H - y^L) + \delta \left[p(1 - 2\mu_t)V(1) + (1 - p(1 - \mu_t))V^- - (1 - p\mu_t)V^+ \right]$ and the cutoff belief is $\mu^{S*} \in (0, \frac{1}{2}]$.

Despite the additional complexity present in the Stochastic setting, our symmetric MPE approximation shares a number of features with the efficient symmetric MPE of the Deterministic setting. There exists a cutoff belief, μ^{S*} , above which players optimally choose REMAIN with certainty. For all beliefs below the cutoff, players mix between REMAIN and SWITCH with an interior probability that depends on the parameters of the model. If $c = 0$, this switching probability places equal weight on REMAIN and SWITCH for all $\mu_t < \mu^{S*}$, while the presence of switching costs, $c > 0$, leads players to switch with increasing (decreasing) probability as their common belief converges towards zero (μ^{S*}). Similar comparative statics with respect to the cost of switching emerge in the Stochastic setting, as well, since increasing c induces a strictly lower probability of switching for all $\mu_t < \mu^{S*}$ and a lower cutoff belief μ^{S*} .

Two main differences distinguish optimal behavior in the Stochastic setting from that of the Deterministic setting. First, since failure does not perfectly reveal incompatibility, the predicted probability of switching in the symmetric MPE approximation necessarily *changes* across periods (with the players' common belief) until they successfully coordinate. This is unlike behavior in the Deterministic setting, where players update their belief to $\mu = 0$ following failure in period 0 and switch with probability $s(0)$ until they achieve success. Second, provided that expected period-benefits for compatible platforms are identical in the two settings and $c > 0$, at a given μ_t , players optimally switch with lower probability in the Stochastic setting, where this relationship is strict for all $\mu_t \in (0, \min\{\mu^{D*}, \mu^{S*}\})$.

B Approximation Method for Stochastic Setting

In this section, we describe our approximation method for computing the equilibrium probability of switching $s^{S*}(\mu_t)$ in the Stochastic setting, where $p \in (0, 1)$. As in the Deterministic setting, we conjecture that there exists a symmetric MPE in which both players choose SWITCH with some probability $s(\mu_t) \in (0, 1)$ for $\mu_t < \mu^{S*}$ and choose REMAIN with certainty for $\mu_t \geq \mu^{S*}$. Recall that the general form of the players' continuation value $V(\mu_t)$ is written as

$$\begin{aligned} V(\mu_t) = & 2s_t(1-s_t) \left\{ p(1-\mu_t) \left(y^H + \delta V(1) \right) + (1-p(1-\mu_t)) \left(y^L + \delta V^- \right) \right\} \\ & + (1-2s_t(1-s_t)) \left\{ p\mu_t \left(y^H + \delta V(1) \right) + (1-p\mu_t) \left(y^L + \delta V^+ \right) \right\} \\ & - s_t c. \end{aligned}$$

For the above candidate equilibrium, this expression can be simplified based on whether μ_t lies above or below the cutoff belief μ^{S*} . That is, for $\mu_t \geq \mu^{S*}$,

$$V(\mu_t) = p\mu_t \left(y^H + \delta V(1) \right) + (1-p\mu_t) \left(y^L + \delta V^+ \right),$$

while for $\mu_t < \mu^{S*}$,

$$V(\mu_t) = \frac{1}{2} \left\{ py^H + (2-p)y^L - c + \delta \left[pV(1) + (1-p(1-\mu_t))V^- + (1-p\mu_t)V^+ \right] \right\}.$$

However, because there are infinitely many possible values of belief μ_t , we are unable to solve the above system of continuation values analytically for a closed-form solution. Moreover, in the absence of a closed-form expression for $s(\mu_t)$, we cannot explicitly calculate the equilibrium cutoff belief μ^{S*} . As a result, we develop a procedure to approximate the probability of switching $s^{S*}(\mu_t)$. We utilize a “guess and check” method with initial cutoff guess $\mu^S = \frac{1}{2}$ and verify that our resulting approximation satisfies the axioms of probability, repeating the process with a new (lower) cutoff guess if necessary.⁴¹ Below are the details of our procedure.

1. **Let** $\mu_0 \in M := (\mu(k))_{k \in \mathbb{Z}}$, **where** $\mu(k) = \frac{(1-p)^k}{1+(1-p)^k}$.

Define M to be the set containing all beliefs of the form $\mu(k) = \frac{(1-p)^k}{1+(1-p)^k}$ for all $k \in \mathbb{Z}$. Clearly, $\mu(k=0) = 1/2 \in M$. Furthermore, if $\mu_t \in M$, then $\mu_{t+1} \in M$ following a failure, *irrespective of the actions chosen by the players*. That is, $\mu^+ \in M$ and $\mu^- \in M$.

This is a direct result of Bayes' rule. In particular, given $\kappa \in \mathbb{Z}$ and $\mu(\kappa) = \frac{(1-p)^\kappa}{1+(1-p)^\kappa}$,

$$\mu^+(\kappa) = \frac{\mu(\kappa)(1-p)}{1-p\mu(\kappa)} = \frac{(1-p)^{\kappa+1}}{1+(1-p)^{\kappa+1}}$$

and

$$\mu^-(\kappa) = \frac{(1-\mu(\kappa))(1-p)}{1-p(1-\mu(\kappa))} = \frac{(1-p)^{1-\kappa}}{1+(1-p)^{1-\kappa}}.$$

Let $\mu_0 \in M$, which guarantees that either $\mu_t = 0$, $\mu_t = 1$, or $\mu_t \in M$ for all $t > 0$.⁴²

⁴¹Our initial guess $\mu^S = \frac{1}{2}$ is informed by the players' optimality condition. Since $\Lambda(\frac{1}{2}) = 0$, SWITCH is strictly dominated by REMAIN at this belief for $c > 0$.

⁴²In our experiment, we induce $\mu_0 = 1/3$, which is $\mu(k=1)$ for $p = 0.5$.

2. **Fix $\mu(k^*) \in M$ such that $\mu^+(k^*) \approx 0$.**

Fixing a belief threshold $\underline{\mu} > 0$, we can solve for $\underline{k}$ that satisfies $\underline{\mu} = \frac{(1-p)^{\underline{k}}}{1+(1-p)^{\underline{k}}}$:

$$\underline{\mu} = \frac{(1-p)^{\underline{k}}}{1+(1-p)^{\underline{k}}} \iff \underline{k} = \frac{\ln\left(\frac{\underline{\mu}}{1-\underline{\mu}}\right)}{\ln(1-p)}$$

This gives us the number of consecutive failures (y^L) following $a_i = a_j$ required for $\mu = 1/2$ to deteriorate to *exactly* $\underline{\mu}$. In general, $\underline{k}$ is not an integer, so define $k^* := \lceil \underline{k} \rceil$, using the ceiling function. For $\underline{\mu}$ sufficiently close to 0, it follows that $\mu^+(k^*) \approx 0$. Thus, in our procedure, we designate $\mu(k^*)$ as the belief that goes to $\mu = 0$ following $a_i = a_j$ and y^L . Note that $\mu^-(k^*) = \mu(1 - k^*) = \frac{(1-p)^{1-k^*}}{1+(1-p)^{1-k^*}}$.

3. **Equipped with initial guess μ^S , construct set $(V(\mu(k)))_{k=1-k^*}^{k^*}$ and solve as a system of simultaneous linear equations.**

Recall that our candidate equilibrium prescribes choosing SWITCH with some probability $s(\mu_t) \in (0, 1)$ for all $\mu_t < \mu^S$ and choosing REMAIN with certainty for all $\mu_t \geq \mu^S$. Given μ^S , we construct the continuation value $V(\mu(k))$ for each $\mu(k)$, $k = 1 - k^*, \dots, k^*$, by imposing our candidate equilibrium structure. The set of continuation values $(V(\mu(k)))_{k=1-k^*}^{k^*}$ is a system of $2k^*$ simultaneous linear equations in $2k^*$ unknowns that can be easily solved.

4. **Compute $s(0)$, $s(1)$, and $(s(\mu(k)))_{k=1-k^*}^{k^*}$.**

As noted in Appendix A.5, for $\mu_t < \mu^S$, the candidate equilibrium probability of switching is


$$s(\mu_t) = \frac{1}{2} \left[1 - \frac{c}{\Lambda(\mu_t)} \right], \quad [\text{B.1}]$$

where $\Lambda(\mu_t) = p(1-2\mu_t)(y^H - y^L) + \delta \left[p(1-2\mu_t)V(1) + (1-p(1-\mu_t))V^- - (1-p\mu_t)V^+ \right]$. For $\mu_t \geq \mu^S$, the candidate equilibrium prescribes play of REMAIN with certainty, $s(\mu_t) = 0$. Using $(V(\mu(k)))_{k=1-k^*}^{k^*}$ derived from Step 3 and equation (B.1), along with $V(0)$ and $V(1)$, we compute the corresponding set of $2k^* + 2$ switching probabilities.


5. **Verify that $s(0)$, $s(1)$, and each $s(\mu(k))$ lie between 0 and 1.**

As a final step, we verify that $s(0)$, $s(1)$, and each $s(\mu(k))$ satisfy the axioms of probability, specifically that each switching probability lies between 0 and 1. If $s(\mu(k)) < 0$ for any k , we consider the largest such k and call it $\hat{k}$. We then set the proposed equilibrium cutoff belief $\mu^S = \mu(\hat{k})$ and repeat Steps 3–5. If $s(\mu(k)) \in [0, 1]$ for all k , we have our desired cutoff belief, $\mu^{S*} = \mu^S$. The resulting set of $s(0)$, $s(1)$, and $(s(\mu(k)))_{k=1-k^*}^{k^*}$ constitutes our symmetric MPE approximation in the Stochastic setting.


Figure B.1 displays the impact of the choice of $\underline{\mu}$ on our prediction. As $\underline{\mu}$ approaches zero, the predicted set of beliefs grows larger in size, approximating the true (infinite) set of beliefs in the limit, given prior $\mu_0 \in M$. Here, further decreasing $\underline{\mu}$ results in negligible changes to the corresponding set of switching rules, which can be seen by comparing Figures B.1a and B.1b. When $\underline{\mu}$ is relatively far from zero, the predicted set of beliefs is small, and the corresponding set of switching rules differs more substantially, as shown in Figure B.1c.



(a) $\underline{\mu} = 0.01$



(b) $\underline{\mu} = 0.05$



(c) $\underline{\mu} = 0.20$

Figure B.1. Predicted equilibrium switching rules by (common) belief, μ_t , in the Stochastic setting.

C Analysis of Switching Rules and CRT Responses

In this appendix, we examine whether individual characteristics are correlated with switching rule decisions, specifically when the players are on known compatible platforms ($\mu_t = 1$). To do so, we construct two alternative measures of cognitive ability (or cognitive process) based on subjects' responses to the questions on the Cognitive Reflection Test (CRT). The first approach we pursue relies on evidence of correlation between CRT score (number of correct responses) and intelligence or cognitive ability (see, e.g., ???). Figure C.1 provides histograms for the number of correct responses out of the seven CRT questions used in our experiment.

Baseline classification. For a baseline classification of cognitive ability, we separate subjects into two types: those with a high CRT score (4–7 correct responses) and those with a low CRT score (0–3 correct responses). In Figure C.2, we compare the mean switching rules for subjects of each type, in each Deterministic treatment, when they are on known compatible platforms (left panel) and provide a scatterplot of the subject-level means (right panel). Both panels illustrate a clear pattern in D10 and D30; subjects with high cognitive ability (measured as a higher score on the CRT questions) chose substantially lower switching rules while on compatible platforms than did subjects with lower scores. In D0, there is no apparent difference based on CRT score, although the relatively smaller number of observations may be influenced partly by one of the outliers.

The same patterns are observed in the Stochastic treatments, as illustrated by Figure C.3, with significantly higher mean switching rules for subjects with lower CRT scores in S0, S10, and S30.

Result 6. *Subjects with a higher score on the CRT questions (CRT-High) choose significantly lower switching rules when on known compatible platforms in all treatments except D0.*

Alternative classification. We also consider an alternative classification, by which we classify subjects into one of three types, corresponding to the types of answers they provided to the CRT questions. We classify a subject as *Reflective* if they provide at least four correct answers (thereby overcoming the instinctive, or intuitive incorrect response) and as *Impulsive* if they provide at least four intuitive incorrect answers, with the remaining subjects classified as *Other*. Across all three Deterministic treatments, the fraction of observations coming from each type is quite similar (no less than 25.9% and no more than 44.8% of the observations for any one category in a given treatment).⁴³ With this classification, the mean switching rules for *Reflective* types are significantly lower than for *Impulsive* types in D10 and D30 (see Figure C.4).⁴⁴ In the Stochastic treatments, the comparison is even stronger. As shown in Figure C.5, in all three Stochastic treatments, subjects classified as *Reflective* select substantially and significantly lower mean switching rules than those classified as *Impulsive*. Altogether, we find consistent correlations between subjects' responses to the CRT questions and their propensity to choose positive switching rules even though they are on commonly known compatible platforms.

⁴³Furthermore, the distribution of types is as follows, with R denoting Reflective, I denoting Impulsive, and O denoting Other: (12R, 8I, 8O) in D0, (15R, 14I, 19O) in D10, and (21R, 19I, 14O) in D30.

⁴⁴Interestingly, the opposite is true in the D0 treatment, although in this case, the mean switching rule by players classified as *Other* is by far the highest.

Figures/hist_crt_8bins-eps-converted-to.pdf

Figure C.1. Histograms of CRT score (number of correct responses out of seven questions), by treatment.

Result 7. *Reflective types choose significantly lower switching rules than Impulsive types when on known compatible platforms in all treatments except D0.*




Figure C.2. Switching rules on known compatible platforms, separated by subjects' cognitive ability (CRT-High vs. CRT-Low)—Deterministic Treatments.




Figure C.3. Switching rules on known compatible platforms, separated by subjects' cognitive ability (CRT-High vs. CRT-Low)—Stochastic Treatments.

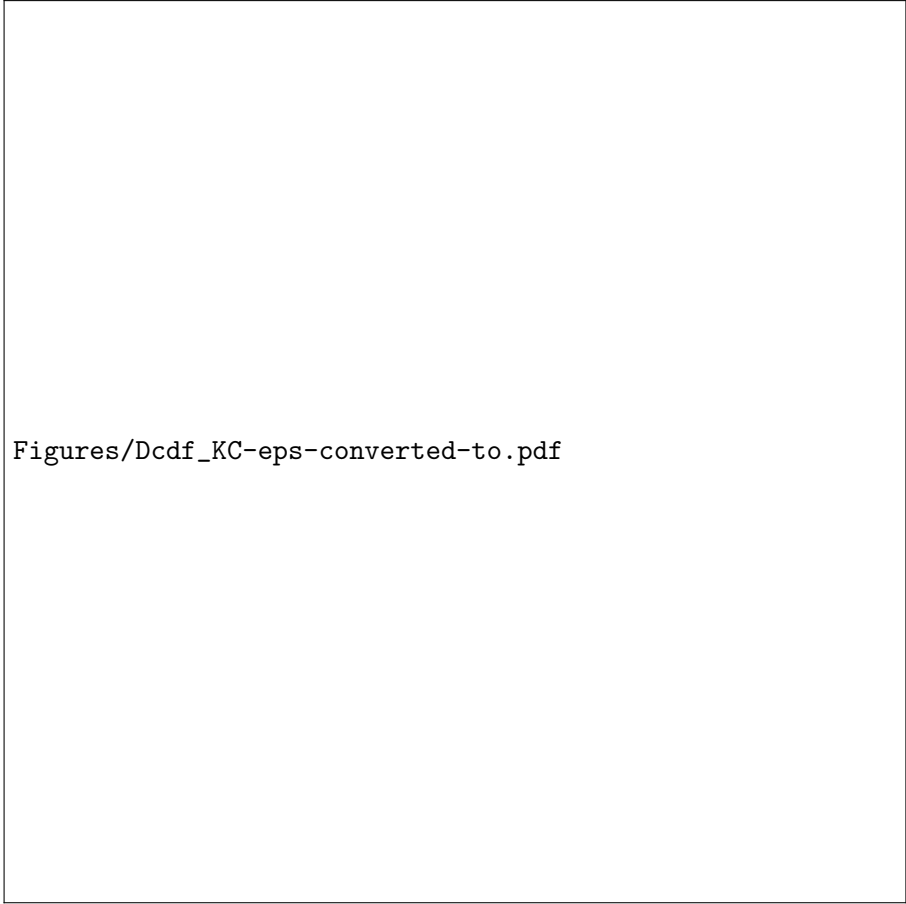


Figure C.4. Mean switching rule by CRT-7 type (*Reflective*, *Impulsive*, or *Other*) in the D treatments, when on known compatible platforms ($\mu_t = 1$).




Figure C.5. Mean switching rule by CRT-7 type (*Reflective*, *Impulsive*, or *Other*) in the S treatments, when on known compatible platforms ($\mu_t = 1$).

D Additional Figures



Figures/Dcdf_KC-eps-converted-to.pdf

(a) Deterministic



Figures/Scdf_KC-eps-converted-to.pdf

(b) Stochastic

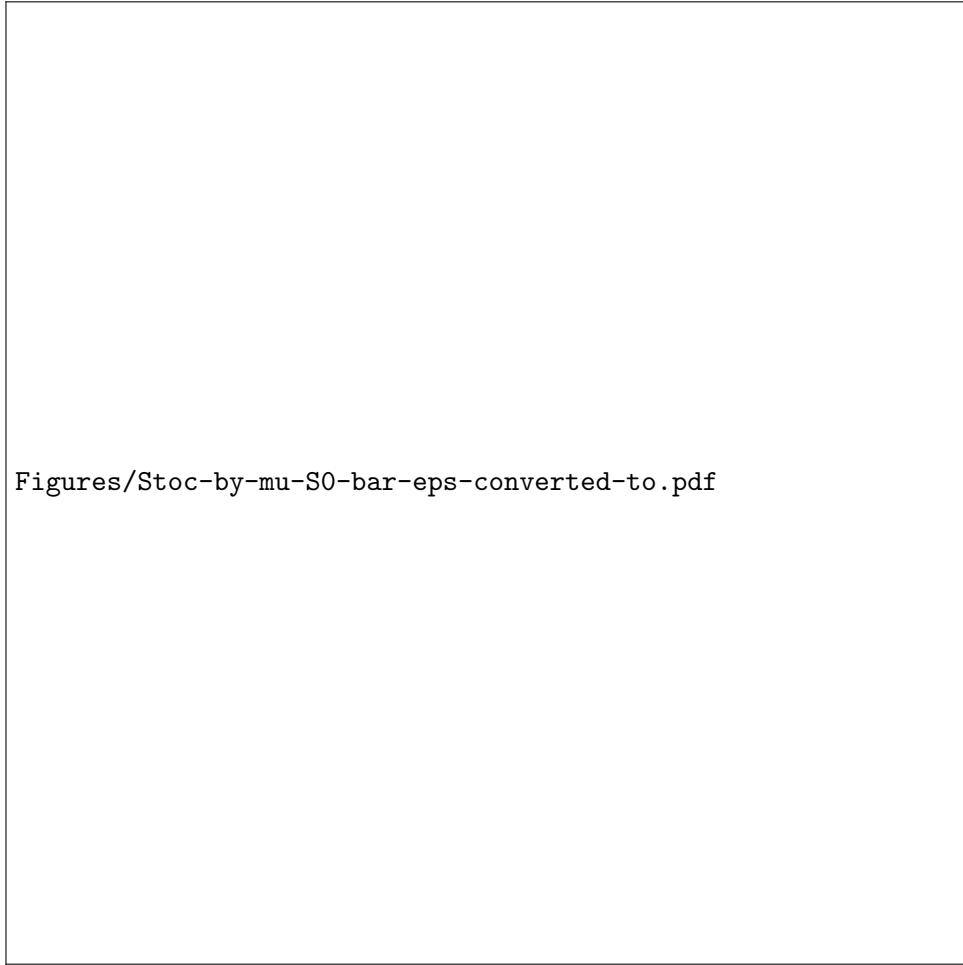


Figure D.2. Mean switching rule by (common) belief, μ_t , in $S0$.

Notes: Error bars indicate standard error of the mean. Dashed (red) lines indicate equilibrium point predictions.




Figure D.3. Switching rules chosen by players on (pooled) likely incompatible versus (pooled) likely compatible platforms, in S0.

Notes: Panel (a): Error bars indicate standard error of the mean.

E Additional Experimental Tasks

E.1 Allocation Game

In Part 1 of our experiment, we implemented the Allocation Game of ?. Table E.1 shows payoffs for each of the fifteen decision rounds in the Allocation Game. All participants completed Rounds 1-15 in the order shown below, making each decision between Option 1 and Option 2 in the role of a Participant X. At the end of the experiment, participants were informed of their randomly assigned role (either Participant X or Participant Y), the round randomly chosen for payment, the option selected by the Participant X in their pair for that round, and their resulting earnings from that decision. Payoffs were converted from points to U.S. dollars at the rate of 60 points = \$1.

Table E.1. Allocation Game (?)

Round	Option 1	Option 2
1	(30,100)	(100,30)
2	(0,0)	(40,30)
3	(30,30)	(60,60)
4	(80,60)	(70,100)
5	(100,30)	(30,230)
6	(70,60)	(90,80)
7	(100,30)	(300,230)
8	(70,60)	(90,50)
9	(30,30)	(30,300)
10	(60,50)	(90,40)
11	(100,30)	(30,230)
12	(230,30)	(100,230)
13	(60,60)	(50,20)
14	(230,230)	(300,230)
15	(60,70)	(50,20)

Notes: Payoffs are denoted in points, with each pair of the form (Participant X, Participant Y).

E.2 CRT Questions

Part 2 of our experiment consisted of seven CRT-style questions.⁴⁵ These questions appeared in the order shown below for all participants, the time limit was 30 seconds per question, and participants earned \$0.80 per correct answer. Participants were informed of their earnings from Part 2 only at the end of the experiment.

1. A man buys a pig for \$60, sells it for \$70, buys it back for \$80, and sells it finally for \$90. How much money has he made? _____ dollars
– Intuitive/incorrect answer: 10, Correct answer: 20
2. A farmer makes 4 piles of hay in one corner of a field and 5 other piles in another corner. If he merges them, how many piles will he have?
– Intuitive/incorrect answer: 9, Correct answer: 1
3. Emily’s father has three daughters. The first two are named April and May. What is the third daughter’s name?
– Intuitive/incorrect answer: June, Correct answer: Emily
4. If you’re running a race and you pass the person in second place, what place are you in? (Please enter 1, 2, 3, or 4)
– Intuitive/incorrect answer: 1, Correct answer: 2
5. In a lake, there is a patch of lily pads. Every day, the patch doubles in size. If it takes 48 days to cover the entire lake, how long would it take for the patch to cover half of the lake? _____ days
– Intuitive/incorrect answer: 24, Correct answer: 47
6. If it takes 5 machines 5 minutes to make 5 widgets, how long would it take 100 machines to make 100 widgets? _____ minutes
– Intuitive/incorrect answer: 100, Correct answer: 5
7. A bat and a ball cost \$1.10 in total. The bat costs \$1.00 more than the ball. How much does the ball cost? _____ cents
– Intuitive/incorrect answer: 10, Correct answer: 5

⁴⁵All of the questions were taken from previous CRT studies: Question 1 from ?, Question 2 from ?, Questions 3 and 4 from ?, and Questions 5–7 from ?.

F Sample Experimental Instructions (S30)

Thank you for participating in today's experiment. I will read through the instructions for each screen to ensure that everyone receives the same information. Please use the chat window in Zoom to ask any questions.

For participating and completing today's experiment, you will receive the participation fee of \$10. In addition, during the experiment, you will have the opportunity to earn more money. Your additional earnings will depend on your decision and the decisions of other participants. At the end of the experiment, you will have the option to request payment by Venmo or by check (which will be mailed to you after you provide a mailing address). No other participant will be informed about your payment.

Please make certain that you have a stable internet connection before beginning today's experiment. It will be very difficult to complete the experiment if you have a spotty connection, and you will disrupt the experiment for everyone else. **If your connection drops during the session, please contact us at (XXX) XXX-XXXX for help with reconnecting. You should write this phone number down now, so that you have it handy if you are disconnected.**

Please do not engage in disruptive behavior during the experiment. If you engage in disruptive behavior, you will receive a warning and, upon a second incident, be removed from the session. If you are removed, you will forfeit the \$10 participation fee as well as earnings from the remainder of the experiment.

Specifically, please pay attention and make decisions in a timely fashion; please do not try to communicate with other participants during the experiment; please do not use your smartphone (or similar electronic device) during the experiment except to call us with questions; and please do not open any additional windows or apps during the experiment.

There are multiple parts in the experiment. Part 1 will be described next. Other parts will be described after you have completed Part 1.

Part 1 (Allocation Game) Instructions

This part is independent of the other parts. It consists of 15 rounds of decisions. In this part, there are two roles: participant X and participant Y. There are an equal number of participants X and participants Y.

At the beginning of each round, the program randomly pairs up each participant X with a new participant Y. The payoff for each participant in a pair is determined solely by the decision of participant X.

Before the first round, the program will randomly assign you to one of the two roles **for the entire part**. However, you will only be informed of your role assignment at the end of the experiment. Instead, **you will all make decisions in the role of a participant X**.

If you learn at the end of the session that the program has assigned you to the role of a participant Y, none of the decisions you have made will count for this part. Your decisions will count only if the program has assigned you to the role of a participant X.

In each round, in the role of participant X, you have to choose between two payoff options for you and for participant Y. For instance, an example of the choices presented to you (different from the actual choices) is shown below.

	Option 1			Option 2	
	X	Y		X	Y
Payoff	50	100	Payoff	100	80

In this example, Option 1 pays you (as participant X) 50 points and pays 100 points to participant Y, while Option 2 pays you 100 points and pays 80 points to participant Y. In each round, you will make the same type of decision between two options.

Once you have made your decisions in all rounds, one round will be randomly drawn for payment by the program. You will be informed about which round is drawn for payment at the end of the experiment. Then, based on your actual role assignment, your decision (if you are participant X) or the decision made by the other participant with whom you are paired in this round (if you are participant Y) will determine your payoff for this part. Points will be converted into US dollars at the exchange rate of 60 points = \$1.

Part 2 (CRT Questions) Instructions

This part is independent of the other parts. In this part, try your best to answer each of the following 7 questions. You have 30 seconds for each question and you will earn \$0.80 for every correct answer at the end of the experiment.

If the answer is a number, please type the number, for example “7” instead of “seven.”

Part 3 (Coordination Game, S30) Instructions

Screen 1. Here are the instructions for Part 3. Part 3 consists of a series of 5 separate matches. In each match, you are randomly and anonymously paired with 1 other participant in this session. Pairs are randomly redrawn between matches. **One of these matches will be randomly selected for payment at the end of the experiment - each match being equally likely.**

The diagram here shows the structure of Part 3. First, pairs are drawn randomly, then you will play Match 1. After completing Match 1, pairs will be redrawn, again at random, and you will play Match 2. Before each new match, pairs will be randomly redrawn.

At no point will you have any information about the identity of the person you are paired with.

Figure F.1. Part 3 Instructions, Screen 1.

Screen 2. So, a match is between you and one anonymous other person. Each of you has two possible color platforms - you can see on the diagram in front of you, one person’s color platform will be either green or purple, while the other person’s color platform will be either orange or blue.

Figure F.2. Part 3 Instructions, Screen 2.

Screen 3. Different combinations of these platforms generate different benefits. There are two cases, the first of which is summarized in table form below. So we will call this Case 1.

In Case 1, the combination of a green platform for one person and a blue platform for the other generates a 50% chance of a benefit of 260 for each person and a 50% chance of a benefit of 120 for each. Likewise, a combination of purple for one person and orange for the other also generates a 50% chance of a benefit of 260 for each person and a 50% chance of a benefit of 120 for each. The other combinations, green with orange, or purple with blue, always generate the lower benefit of 120 for each person.

In this case, we refer to green and blue as compatible, and we refer to purple and orange as compatible, since those pairings generate a 50% chance of the higher benefit of 260 for each person.

Figure F.3. Part 3 Instructions, Screen 3.

Screen 4. In the second case, shown on the right, the combinations of color platforms that are compatible have switched. That is, in this case, a combination of green for one person and orange for the other person generates a 50% chance of a benefit of 260 for each person and a 50% chance of a benefit of 120 for each. Likewise, for a combination of purple for one person and blue for the other. The remaining combinations, which are green with blue, or purple with orange, always generate the lower benefit of 120 for each person.

So, in this second case, we refer to green and orange as compatible, and we refer to purple and blue as compatible.

These are the two possible cases. Before each match, the program will randomly select one of these, Case 1 or Case 2, for the entire match. Thus, the selected case will not change between periods within a match, but may change between matches. Throughout the experiment, neither you nor the other person will ever be told which case was selected.

Figure F.4. Part 3 Instructions, Screen 4.

Screen 5. Each match consists of a sequence of periods. Before period 1, you and the other person are each assigned an initial color platform. Note that only you can observe your own initial platform. The initial platforms assigned to you and the other person have a 1 in 3 (33.33%) chance of being compatible.

Suppose that the compatible platforms, as shown on the left, are green with blue, and purple with orange. **Reminder:** in the experiment, you won't be told that these are the color platforms that are compatible. However, in order to illustrate how the initial platforms are determined, you can see this information for the current example. Moreover, throughout the instructions, any information that is not observed during the experiment will be shown inside brackets, as it is here.

So, **for this example**, if you are assigned an initial platform of green, then there is a one in three chance that the other person is assigned an initial platform of blue (which is the

compatible platform for green), and a 2 in 3 chance the other person is assigned an initial platform of orange.

If instead you are assigned an initial platform of purple, then there is a 1 in 3 chance that the other person is assigned an initial platform of orange (which is the compatible platform for purple), and a 2 in 3 chance the other person is assigned an initial platform of blue.

Once again, in the experiment, you will not be told which platforms are compatible and you will not be told the other person's platform at any point.

Figure F.5. Part 3 Instructions, Screen 5.

Screen 6. To make this clear, on this screen, we've replaced the other person's color platforms with question marks in the table and updated what you can infer based on your own initial platform. If you are assigned to green, then there is a 1 in 3 chance that the other person is assigned to the color platform compatible with green, and a 2 in 3 chance that they are not. If you are assigned to purple, then there is a 1 in 3 chance that the other person is assigned to the color platform compatible with purple, and a 2 in 3 chance that they are not.

Figure F.6. Part 3 Instructions, Screen 6.

Screen 7. In each period, you will decide on an action. Your decision will determine whether you remain on your current platform or switch to the other platform. You can see the two options here. Remain, which is shaded in white, has a cost equal to 0. Switch, which is shaded in yellow, has a cost equal to 30 points.

Your Payoff in each period is equal to Your Benefit minus Your Cost for that period.

In the diagram below, you can see a timeline of how the match progresses. At the beginning of the match, each of you will be assigned initial platforms. Only you observe your own initial platform. You then make decisions that determine your actions, and based on your initial platforms, your actions will determine updated platforms. That is, if your action is Remain, your updated platform will be the same as your initial platform, while if your action is Switch, your updated platform will switch from your initial platform to the other color platform.

Once the platforms are updated, your payoff will be determined, based on the benefit generated by the updated platforms and the cost of your own action. At the end of the period, the updated platforms will become the current platforms at the beginning of the next period.

In period 2, and in every subsequent period, everything progresses in the same way, with the updated platforms at the end of a period becoming the current platforms at the beginning of the next period.

Figure F.7. Part 3 Instructions, Screen 7.

Screen 8. To help illustrate further, we're now going to explain some payoff examples. We'll begin by setting up a scenario in place at the beginning of a period. Once again, things that are shown in brackets represent things that you will not be told during the experiment. However, we show them to you for the purposes of the example so that you can understand the way payoffs are determined.

So consider the scenario shown on the left side of the screen. Suppose you are assigned to the initial platform of purple, or that your current platform (in periods after the first) is purple. Likewise, suppose the other person's initial or current platform is orange. Remember, you will not be told that this is their platform.

In addition, suppose that the program determined that the compatible platforms for this match are purple with blue and green with orange. Again, you would not know that these are the compatible platforms during the experiment. The resulting benefits generated by different combinations of platforms are shown in the table.

Figure F.8. Part 3 Instructions, Screen 8.

Screen 9. Consider Example 1 here. Suppose your action is remain, and the other person's action is remain. Note that each person observes both actions, but only after both decisions have been made.

Then your own updated platform will remain purple, the same as your current platform. The other person's updated platform will remain orange. So while you would be told that the other person's action was remain, you would still not be told which color platform they remain on.

Based on the updated platforms, the program would determine the benefits from the table in the bottom left. For this combination of updated platforms, purple and orange, which is outlined in red, the benefits are 120 points each. That is, purple and orange are not compatible.

Furthermore, since each person's action was remain, each person pays a cost equal to 0. Then, your payoff in this period would be your benefit, 120, minus your cost, 0, for a period payoff of 120. Likewise, the other person's payoff would be the benefit, 120, minus their cost, 0, for a period payoff of 120 as well.

Figure F.9. Part 3 Instructions, Screen 9.

Screen 10. Next, consider Example 2. The setup shown on the left hand side is the same as for the last example. In this case, suppose your action is switch, while the other person's action is remain. Again, these actions are shown to both you and the other person.

Your updated platform will switch from purple to green. The other person's updated platform will remain orange.

Since these are compatible platforms, the benefit has a 50% chance of being 260 points for each of you, and a 50% chance of being 120 points for each of you.

Your payoff in this period would be your benefit minus your cost, which is 30 points in this case, since your action was switch. The other person's payoff would be the benefit minus their

cost, which is 0 in this case, since their action was remain. Thus, there is a 50% chance that your payoff is 260 minus 30, or 230 points, and the other person's payoff is 260 minus 0, or 260 points; and there is a 50% chance that your payoff is 120 minus 30, or 90 points, and the other person's payoff is 120 - 0, or 120 points.

Note that the benefit from compatible platforms need not be the same in different periods. That is, in each period, compatible platforms have the same 50% chance of 260 each and 50% chance of 120 each; you can think of this in terms of coin flips – even if the coin flip generates a benefit of 260 each in one period, if the platforms in the next period are still compatible, there is a new coin flip which may result in 260 for each or 120 for each.

Figure F.10. Part 3 Instructions, Screen 10.

Screen 11. Finally, consider Example 3. Again, the setup shown on the left hand side is the same as for the other two examples. In this case, suppose your action is switch, and the other person's action is switch.

Then your own updated platform will switch from purple to green. The other person's updated platform will switch from orange to blue.

The updated platforms, green and blue, are not compatible, so the benefits would be 120 points for each person (outlined in red). Furthermore, since each person's action was switch, each person pays a cost equal to 30 points. Thus, your payoff in this period would be your benefit, 120, minus your cost, 30, for a period payoff of 90 points. Likewise, the other person's payoff would be the benefit, 120, minus their cost, 30, for a period payoff of 90 points as well.

On the next page, you'll see a few quiz questions to answer regarding the calculation of payoffs in another example similar to these.

Figure F.11. Part 3 Instructions, Screen 11.

Screen 12 (Quiz 1). For these questions, you can refer to the setup shown on the left side of the screen. That is, suppose your current platform is blue. The other person's current platform is green.

In the bottom left, you can see that the program selected the case where blue and purple are compatible and where green and orange are compatible.

Figure F.12. Part 3 Instructions, Screen 12.

Screen 13 (Quiz 1 Answers). Let's quickly review the answers to each question. First, during the experiment, you will not be told the other person's platform.

Second, since you started on Blue and your action was Remain, your updated platform is also Blue. The other person started on Green and their action was Switch, so their updated platform would be Purple.

Then, based on the updated platforms, the program would determine the benefit. From the table, it's a 50% chance of 260 each and a 50% chance of 120 each. Since your action was Remain, you do not pay a cost, so the benefit is also your payoff.

For the other person, the same applies for the benefit, but since their action was Switch, they pay the cost of 30 points, which is why their payoff is a 50% chance of 230 points and a 50% chance of 90 points.

Figure F.13. Part 3 Instructions, Screen 13.

Screen 14. Next, let's explain how you make a decision. On the screen you can see a partial screenshot of what you will be shown when you are asked to make a decision. On the left hand side, you can see a table describing the possible benefits with the other person's possible platforms replaced by question marks.

You make your decision on the right part of the screen. You can do so in a couple of ways. First, if you want to SWITCH, you can click on the SWITCH button (highlighted in yellow). After you click on the SWITCH button, all of the numbers in the 10 by 10 grid on the far right will be highlighted in yellow. Please go ahead and click the yellow SWITCH button.

If you want to REMAIN on your current platform, you can click on the REMAIN button (shaded white). After doing so, the numbers in the 10 by 10 grid will be shaded in white. Go ahead and click the white REMAIN button.

You can also randomize your action using what's called a switching rule. That is, you can direct the program to choose SWITCH with some chance and to choose REMAIN with any leftover chance. In particular, the switching rule works as follows:

First, type in an integer number from 0 to 100 in the entry box and click the gray button labeled "Update". In the 10 by 10 grid, all of the boxes with numbers less than or equal to the integer you enter will be highlighted in yellow. All of the boxes with numbers greater than the integer you enter will remain shaded in white. You can try this out for yourself on the screen, by entering a number and clicking Update.

A decision is finalized by clicking the red "Confirm" button. After you confirm, the program will draw a random integer between 1 and 100. If the random integer is equal to or lower than your switching rule (which means it will be highlighted in yellow in the grid), then your action will be SWITCH and you will pay the cost. If the random integer is higher than your switching rule (meaning it is shaded white in the grid), then your action will be REMAIN and you will not pay the cost.

Thus, the switching rule allows you to specify the number of chances out of 100 with which you want your action to be SWITCH (and therefore also the leftover number of chances out of 100 with which you want your action to be REMAIN).

Before you click the Confirm button, make sure that the highlighted boxes in the 10 by 10 grid match your decision. If they do not, you may need to click Update, or click on the desired Switch or Remain button, before clicking Confirm again.

Please note that you should make and CONFIRM your decision within 30 seconds in order for the experiment to proceed in a timely fashion. Thus, please pay attention to the screens and make sure to click the Confirm button when you are done.

Figure F.14. Part 3 Instructions, Screen 14.

Screen 15. Here is an example of how the switching rule works. Suppose you select a switching rule of 73. This randomizes your decision so that your action has a 73% chance (i.e., 73 chances out of 100) of being SWITCH, and a 27% chance (i.e., 27 chances out of 100) of being REMAIN.

Given this example, if the random integer drawn by the program were 43, your action would be SWITCH (and you would pay the cost), since 43 is less than 73 (and so would be shaded yellow). If the random integer drawn by the program were 76, your action would be REMAIN (and you would not pay any cost), since 76 is greater than 73 (and would be shaded white).

If you haven't already, you can enter your own example Switching Rule and see the grid update, as well as an explanation of the chances for each action under that rule. Make sure to click the gray update button after entering a switching rule so that the grid updates before clicking Confirm.

Finally, note that clicking the SWITCH button is exactly the same as entering a switching rule equal to 100, since both will lead to the action SWITCH for certain. Likewise, clicking the REMAIN button is the same as entering a switching rule equal to 0, since both will lead to the action REMAIN for certain.

Figure F.15. Part 3 Instructions, Screen 15.

Screen 16. Next, we turn to the length and timing of a match.

The number of periods in each match is randomly determined, according to the following procedure. At the end of each period, there is a 9 out of 10 or 90% chance that the match will continue for at least another period. However, if the match ends before 10 periods have been completed, (i) you will not be told right away, and (ii) you will continue making decisions through period 10.

For example, take a look at the Match length example 1 at the bottom of the screen. Each circle represents a period, numbered from 1 to 10. In this example, the match ended after period 5; thus, the first 5 circles are white, and the circles from 6 to 10 are shaded gray. That is, after each period, there was a 9 out of 10 chance that the match would continue, and it did up until after period 5, when it ended. Nevertheless, even though the match ended after period 5, you would not be told right away, and you would continue to make decisions through period 10.

After period 10, you will be told whether the match has already ended and, if so, after which period. Your total payoff for the match is the sum of payoffs from all periods before the match ends. Thus, payoffs from any decisions made after the match ended will not count.

So, in the example at the bottom of the screen, you would not be told until after period 10 that the match ended after period 5. Furthermore, your total payoff for the match would be the sum of your payoffs from periods 1 through 5 only. That is, periods 6 through 10 would not be counted towards your match payoffs.

If the match has not ended after 10 periods, you will be told and then you will continue making decisions 1 period at a time, until the match ends.

Figure F.16. Part 3 Instructions, Screen 16.

Screen 17. To see this, consider the match length example 2 at the bottom of this screen. In this example, the match did not end until after period 12. Thus, in the first 10 periods, after each period, there was a 9 out of 10 chance the match would continue, and each time, it did. After period 10, you would be told that the match had not yet ended, and you would continue on to period 11.

After period 11, there was again a 9 out of 10 chance the match would continue and it did. You would be told that the match did not end, and you would continue on to period 12. After period 12, there was again a 9 out of 10 chance the match would continue, but it did not. You would be told that the match ended, and no more decisions would be made.

In this example, your total payoff for the match would be the sum of your payoffs from all 12 periods.

Figure F.17. Part 3 Instructions, Screen 17.

Screen 18. On this screen, you can see some sample match lengths that were pre-drawn according to the procedure described on the previous page. Using a computer to roll a virtual 10-sided die, the match length equals the number of rolls it takes to roll a 10. If the die roll is 1 through 9, the computer would roll again. Once a 10 is rolled, the count stops and the number of rolls is the match length.

We did this 50 times and plotted a bar graph to show you a sample sequence of the resulting match lengths.

For perspective, in large samples of these match lengths, the expected average match length, given the 90% chance of continuing after each period, is equal to 10.

Figure F.18. Part 3 Instructions, Screen 18.

Screen 19 (Quiz 2). On this screen there are a couple of questions about the number of periods and the timing of a match. Please answer each question.

Figure F.19. Part 3 Instructions, Screen 19.

Screen 20 (Quiz 2 Answers). Again, let's briefly review the answers. For the first part, even though the match ended after period 7, you would not find out until after period 10. You will make decisions in 10 periods. However, only your payoffs from the first 7 periods will count towards the payoff for the match.

For the second case, the match ended after period 15. You would make decisions in 15 periods. You would find out after period 15 that the match ended and your total payoff will be

the sum of the payoffs from all 15 periods.

Figure F.20. Part 3 Instructions, Screen 20.

Screen 21. Next we have an example screenshot of the feedback you will see at the end of each period, including the history panel shown at the bottom of the screen.

On the left side, you can see the table summarizing the possible benefits alongside your two possible platforms. On the right side is the feedback that you will see after decisions have been made.

You will be shown Your Action - in this case, it was SWITCH - and Your Updated Platform - which in this case was Orange. You will also be shown the other person's action - in this case, it was REMAIN - but not their updated platform.

Note also that the feedback only shows the actions, but does not report any switching rules that were used by either person. Thus, if you used a switching rule to determine your action, the other person will not know, and if the other person used a switching rule to determine their action, you will not know.

In the example shown here, Your Benefit is reported as 260. Given your updated platform is Orange, this allows you to infer that the other person's updated platform is the one that is compatible with Orange, since that is the only way to obtain the higher benefit of 260 if you are on Orange. If, instead, the reported benefit were 120, you might not necessarily know if the other person's updated platform is compatible with Orange or not, since the lower benefit of 120 is possible in either case.

For this example, Your Benefit is 260, Your Cost is 30, since your action was SWITCH, and thus, your payoff for the period is 230 points. All of this information is listed in the history panel at the bottom of the screen. As a match progresses, you will see the history panel throughout all stages of each period.

Below the feedback, there is also a reminder box. Here, during the first 10 periods, you'll be reminded that only after period 10 will you learn whether or not the match has ended. Then, after period 10, this is also where you will see the information about whether the match has ended after each subsequent period, if applicable.


Below that is a reminder is that after each period, if the match has not yet ended, there is a 90% (or 9 out of 10) chance that the match will continue for at least another period. And the third reminder is that you are paired with the same other person for all decisions in the current match.

Figure F.21. Part 3 Instructions, Screen 21.

Screen 22. Here is a final summary before we begin Part 3.

Please note that points from the randomly selected match for your payment will be converted to dollars at the rate 180 points equals \$1. Please take a moment to browse the other bullet points.

Finally, if you are disconnected at any point, please contact us at the number (XXX) XXX-XXXX. Again, if you did not already write it down, please do so now. We ask that you remain attentive throughout the experiment, and please make your decisions in a timely fashion.



thesis-template/Figures/Stochastic-30/22_redacted.png

Figure F.22. Part 3 Instructions, Screen 22.