Mimicking and Modifying: An Experiment in Learning From Others

David Glick
Assistant Professor
Department of Political Science
Boston University
Dmglick@bu.edu

C. Daniel Myers
Robert Wood Johnson Scholars in Health Policy Program
University of Michigan
myerscd@umich.edu

January 20, 2012

Abstract

Actors often make decisions without knowing exactly how their choices produce outcomes. In some instances many actors make a similar decision, providing opportunities for learning from others' choices. Canonical models of learning in this situation rely on a kind of policy uncertainty which assumes that that individuals can learn all information that other actors have simply by observing their actions. However, in many situations, knowing an actor's action and her goal tells one something, but not everything, about the actor's private information. Drawing on a new representation of policy uncertainty which has gained traction in top journals, we develop and test a model of learning in an environment of policy uncertainty in which actors can only learn something from others' actions and outcomes. While this model better approximates many features of learning from others, and has great theoretical appeal, it places greater demands on decision makers and requires them to use non-obvious strategies to optimize. We conduct an experimental test to examine whether subjects can learn to act optimally in this more complicated learning environment. We find that subject behavior deviates in important ways from optimal behavior. These deviations have important implications for these models' predictions and for their use to study learning from one's own, and others' experiences.

1 Introduction

Individuals and organizations regularly make challenging decisions without knowing exactly how their choices produce outcomes. In some instances many will confront similar challenges at roughly the same time providing opportunities to learn from, or follow, others' choices and policies. These opportunities may enable them to achieve outcomes superior to those they would have achieved by acting independently. However, to do so they must extract the proper information from others' choices. Frequently this learning is complicated by the fact that others are in somewhat different situations and have somewhat different goals. Common tasks such as designing a new syllabus and buying wine are examples of these situations. This description is also highly applicable in political and policy decisions in a wide range of contexts and applications. In short, understanding how well people can learn from others' experiences and expertise to make the own choices is an important empirical question with a wide range of applications and implications.

Modeling situations like those above requires formalizing learning in situations with incomplete information. Models where information about choices' outcomes is incomplete are foundational in behavioral and institutional political science. Some of these models also involve opportunities to learn from others' actions. For example, in Gilligan and Krehbiel (Gilligan and Krehbiel, 1989) the floor of a legislative body can learn which policy will deliver its preferred outcome by observing the action of a committee that has a more expertise but a different ideal policy. Similarly, political behavior scholars have long believed that uninformed citizens can learn about politics by examining the actions of better informed members of their social networks (Berelson et al., 1954; Huckfeldt and Sprague, 1995). Finally, in the large diffusion literature, one government's decision to adopt a new policy may be influenced by learning about other governments' policies. In all of these settings a first actor takes an action, whether it is writing a law, adopting a policy position, or choosing a candidate to vote for. A second actor has to make a similar decision with information about the first's action (and possibly outcome) and the difference between the first actor's ideal outcome and her own.

Theoretical results in these situations depend on the particular conception and formalization of policy uncertainty. In the canonical model (Gilligan and Krehbiel, 1989), uncertainty is assumed to be a linear shock that is added to the chose policy to produce an outcome. Further, this linear shock is the same for all policy choices and all actors. Thus, when one actor chooses a policy, her outcome is the sum of that choice and some linear shock. When a second actor chooses a policy, his outcome is his policy choice plus the same linear shock that was added to the first actor's policy. These features of the policy uncertainty mean that any actor who observes the policy choice and outcome of another actor can deduce the size of the linear shock, and thus choose the "perfect" policy. These "fully invertible" signals are the canonical model's biggest shortcoming (Callander, 2008, 2011a,b). A player can "invert" (reverse engineer) the information to extract all of the available knowledge and apply it to their own decisions.

In response to these limitations, Callander (2008) proposed a different sort of policy uncertainty which uses the mathematics of Brownian motion to represent partially invertible signals. The key difference is that rather than a constant linear shift from policies to outcomes, this representation features outcomes drawn from a symmetrical distribution around a known linear trend. Actors know, for example, that on average, a "bigger" policy will produce a "bigger" outcome, but they do not know the exact relationship at other locations in policy space. Moreover, the variance increases with the distance from the known policy so actors know less the further they move policy from the known one. This uncertainty representation thus offers a better representation of the kinds of decisions made by actors seeking to learn from others. Observing an earlier actor teaches something, but not everything, about the policy choice, and teaches more about the effects of policies that are more similar to the observed policy than about policies that are less similar.

To illustrate, consider a governor who wants to increase the percentage of state residents with Medicaid insurance coverage in her state by 10 percent, but is unsure of the effect that increased spending will have on coverage. She observes a neighboring state increase spending by five percent, a change that results in a seven percent increase in coverage. If the assumptions of the canonical model are true, then the first state's governor can increase spending by eight percent and know with certainty that this will increase coverage by ten percent. In this case, the signal from the first governor's action is fully invertible, in that it allows the observer to learn all pertinent information about the policy uncertainty. Of course, the uncertainty surrounding policy decisions is often more complicated than this. Different policy decisions may be subject to different "shocks" such that making slight changes to policies with known outcomes may produce very different outcomes. Increasing spending by eight percent might produce an even greater increase in coverage (if, for example, those gaining coverage are young adults who are cheap to insure) or less of an increase in coverage (perhaps those gaining coverage are hard to reach and require more advertising spending to make them aware of the program). The key point is that the second governor learns a great deal from the first governor about actions that are similar to the first governor's action, but little about actions that are less similar. In this second case, the signal that the second governor receives from the first's action is partially invertible.

We conducted an experiment to test individuals' ability to learn optimally in a richer and more complicated information environment than most existing work assumes. Recently, Callander (2008; 2011a; 2011b) has introduced an intuitive and mathematically compelling representation of policy-outcome uncertainty based on Brownian motion. This formalization elegantly incorporates a number of attractive properties including partially invertible signals such that actors can learn something, but not everything, about others policies from observing a policy-outcome pair, increasing uncertainty with larger moves away from known locations in policy space, and an simple representation of variable decision complexity. We develop a model of "mimicking" and "modifying" based on Callander's work and we test its clear behavioral predictions using individuals in the laboratory. This elaboration adapts the uncertainty representation of Callander's model to the situation of learning from other actors, and is the first test of Brownian motion as a representation of partially invertible signals. Specifically, we offer the first test of a model using the Brownian motion uncertainty representation (Callander, 2008, 2011a,b). While this analytical tool offers great potential as evidenced by recent articles in both the American Political Science Review and the American Economic Review, this potential, especially outside the modeling community, depends on whether humans can incorporate information according to the model's prescriptions.

Our experiment tests the empirical viability of the Brownian motion uncertainty representation in general and tests specific predictions about learning from another's choice. It thus provides insights about Brownian motion's plausibility as an uncertainty representation in a variety of applications in general and about human capacity to rationally learn from others' choices and their outcomes more specifically. These insights contribute in two closely related, but distinct, ways. They contribute as "basic science" with a wide range of applications to problems featuring incomplete information by asking how well people can learn from partially invertible signals and whether they can behave in ways consistent with a new and appealing uncertainty formalization. Secondly, as we test predictions of participants' ability to rationally incorporate information from another's policy choice, it contributes more specifically to literatures concerned with learning from others' choices in institutional and individual settings including policy diffusion, voter learning, and learning from one's own experiments.

We first review the literature concerning models of policy-outcome uncertainty, learning from others, and experimental tests of decision making in these settings. We then begin with Callander's formalization and derive predictions about "mimicking" and "modifying." Next, we describe our experimental investigation. Finally, we report the results. To preview them, our results show that participants can learn, but do so slowly, and are not always able to adopt optimal strategies.

2 Background

Modeling decisions when the relationship between policies and the outcomes the produce is a common challenge. Doing so formally requires specifying a "policy mapping function" by which policies (choices) produce outcomes (Callander, 2008). The workhorse mapping function in continuous space (e.g. Gilligan and Krehbiel, 1987, 1989) is a linear shock ω drawn from a known distribution. In this canonical model, originally applied to delegation and committee expertise in Congress, the outcome (y) that a policy (x) produces is $x + \omega$. As Callander (2008) argues persuasively, this assumption inadequately captures policy complexity in many instances. Its main shortcoming is that it is perfectly invertible. Because the shift ω is constant at all locations in policy space, actors can know the relationship between all policies and outcomes after inferring omega. In some applications actors can infer ω by subtracting x from y after observing a policy and its outcome and thus learn from a previous policy experiment. In other applications, particularly expertise models in which experts are assumed to possess information about the uncertainty, one can subtract an expert's policy choice from her ideal outcome without needing to observe the policy's outcome.

While the details vary by application, and may be more subtle in strategic settings like delegation, the workhorse uncertainty model has its limitations. Without rehashing Callander's (2011b) very thorough critique and explanation, we briefly highlight the intuition. Most generally, the problem is that the factors that make it difficult to tailor a policy to achieve a particular outcome are not the same ones that make it hard to tailor a policy to achieve a different outcome. One state's experiment in a challenging policy area does not always eliminate the uncertainty about policies in other states with different goals. Similarly, one competent committee's markup of bill does not necessarily shed a lot of light on the outcomes that a different bill on the same topic will produce; this is particularly when the bills are complex. Moreover, in many instances, the uncertainty about the connection between polices and outcomes grow the further one moves away from a relatively well understood policy. For example, one may be able to alter a graduate Congressional Politics syllabus for an advanced undergraduate seminar with confidence, but knowing what works and what does not in the graduate class will be relatively unhelpful in designing an introductory lecture class.

Callander proposes an alternative model of the uncertain connection between policies and outcomes, and readdressed the delegation problems to which this model of uncertainty was originally applied. His theoretical and methodological innovation is reformulating policy-outcome uncertainty by modeling the shock as the realization of a Brownian motion. Uncertainty is represented by a random walk with variance variance σ^2 around a known linear trend which is called the "drift" (μ) Actors are assumed to know that the policy process is a Brownian motion and to know its parameters. (See figure 3.1, similar to figure 1 in Callander (2011a), for an example of a Brownian path.) Observing an informed actor's policy choice reveals one point through which the function passes. One thus learns something, but not everything, about other policies' mappings. To help with intuition, consider a boat known to be drifting northeast with the current. If it is also subject to random wind and waves, we would know it moves northeast on average but with variation around that "drift." If we even know one point which it passed through we would know a lot more about its location further along its drift northeast. The further that the boat has moved from the point where we originally sighted it, the less accurate our guess about the boat's location will be. The same is true about different policies in the model.

While seemingly complex, the math behind the Brownian motion formalization is quite simple and offers a number of attractive properties. Given one known policy-outome pair, that is, one known point on the Brownian path, the expected value of the outcome that another policy produces is very simple. The expected outcome is just a linear extrapolation just like in the canonical model but instead of being deterministic it is stochastic. The variance around this expected outcome is the product of the Brownian variance and the size of the move from the known policy to the new one. In other words, the uncertainty about policies increase the more dissimilar they are from "known" policies. Actors can know roughly (in expectation) which direction to move from another's policy to get to their own ideal outcome, and how far they should move. This nicely approximates many real life situations. We often may not know exactly how policies produce outcomes, but we do have a sense of how to tweak an existing policy to achieve different goals. It also approximates many situations in which one can be relatively confident about small changes and relatively uncertain about the consequences of more radical policies. Finally, one can easily model relatively complex or relatively simple choice environments by changing the random walk's variance. The larger the variance, the less confident actors can be about untested policies' outcomes.

Callander (2011a; 2011b) extended his model to investigate when actors will experiment with new policies, and Glick (2010) has adopted this formalization of uncertainty to model learning from others. We seek to test one aspect of these models by examining whether participants extract the proper information from partially invertible Brownian motion signals. In particular, we are interested in how people react to two different challenges posed by Brownian uncertainty: the challenge posed by decreased similarity between their ideal outcome and the outcome produced by the known policy-outcome pair, and the challenged posed by increased policy complexity. Decreased similarity and increased complexity both make it more difficult for people to incorporate information about a known policy-outcome pair into their decision-making. Equilibrium predictions suggest that in both situations policy makers should shade their policy choice closer to the policy whose outcome is known, in some cases simply mimicking the known policy and ignoring information about the distance between the known policy's outcome and their ideal outcome. However, actors may not be able to optimally incorporate information in this complex information environment, and may respond to these two challenges differently. Examining how experimental subjects react to these two challenges will provide valuable information about when actors will optimally learn in complex information environments, and when these models fail to describe human behavior.

Other models of social learning and information transfer, some of which have been tested in the lab, particularly information cascade models (Bikhchandani et al., 1992, 1998), have demonstrated the importance of information invertibility assumptions. In these sequential choice models, signals quickly shift from invertible, to partially invertible, to non-invertible as more and more actors take actions. These models demonstrate how wildly divergent and ostensibly irrational actions can follow from completely rational behavior under the right conditions. While the model we seek to examine relates to this work, it is different in a couple of key ways. For one, it is generally more concerned with decision making and problem solving strategies in which signals are partially invertible whereas some of the other models are primarily concerned with the dynamics of information transfer as signals become less and less invertible. Relatedly, we focus on one actor facing a decision after others have faced similar, but potentially different, decisions. We are less concerned with many actors facing the same choice in sequence.

Our specific theoretical extension of the Brownian motion framework and the experimental test of its behavioral predictions focused on learning from another's choices and outcomes when signals are partially invertible. When extended to institutions, this work contributes to the large and growing diffusion. This broad literature is concerned with "uncoordinated policy interdependence" (Elkins and Simmons, 2005) which is at least partly attributable to governments learning from others' policies. (See Dobbin et al. (2007); Karch (2007) for recent reviews.) Diffusion studies have focused on sub-national policy making (e.g. Grossback et al., 2004; Volden, 2006), amongst nation states, (Meseguer, 2006; Weyland, 2005), and business firms (Strang and Still, 2004). This literature is largely empirical and inductive. Diffusion problems are inherently complex and causal inference is difficult. Much of the literature inadequately distinguishes plausible diffusion mechanisms (Dobbin et al., 2007) from each other. Moreover, as Volden et al. (2008) demonstrate formally, distinguishing diffusion from actors independently adopting similar policies is also difficult. For all of these reasons, deductive formal theory combined with experimental tests of precise behavioral predictions marks a methodological divergence (exceptions include Baybeck et al., 2011; Tyran and Sausgruber, 2005; Volden et al., 2008) and offers the potential to overcome other works' limitations. While the literature has cited a number of diffusion mechanisms and incarnations including emulation, coercion, imitation, competition, and legitimacy (Dobbin et al., 2007; Elkins and Simmons, 2005; Karch, 2007; Shipan and Volden, 2008), we focus specifically on information and learning. Our model splits the difference between fully rational learning models (Meseguer, 2004) in which actors glean all available information from others' experiences and update their beliefs via Bayes' Rule, and boundedly rational learning (Weyland, 2005) in which actors rely on behavioral heuristics (e.g. availability and representativeness).

While the diffusion literature, as well as the Callander and Glick models, largely concern institutions, many of the insights apply to individual behavior as well. The formal analysis treats decision makers as unitary actors and the apparatus is thus transferable to individual decisions. As early as Downs (1957) and Berelson et al. (1954), scholars of political behavior have argued that while most voters are uninformed they follow informed members of their social groups to make political decisions. Lupia and McCubbins (1998) repeat this claim and argue that it can be rational for voters to do so. However, these cue taking literatures generally focus on binary decisions and invertible signals. The model on which we base our test offers a richer conception of the types of decisions and information involved in attempts to learn from others.

A few experiments, in addition to the cascade studies, have examined policy diffusion across actors in a laboratory setting. Tyran and Sausgruber (2005) examine the adoption of tax policies in small laboratory markets and find that individuals learn from each other but fail to optimally update their tax policies in response to this information. Ahn et al. (2009) examine separate models of voters learning from others in their social networks where some players have an incentive to mislead other players. They find that experimental participants tend to trust more informed participants too much. They copy their actions and ignore the possibility that they are being intentionally misled. Again, all of these experiments deal with fully invertible signals.

3 Model Overview: Brownian Motion and Learning

The predictions we tested come from a broader learning model (Glick, 2010) which implements Callander's uncertainty representation to problems of learning from others. We focus only on the "mimicking" and "modifying" tradeoffs for one actor which acts after another. Briefly, an actor, z_i (one could consider many actors z_i: $i \in 1...N$) faces a complicated policy decision. Each policy p_i produces an outcome $o_i \in R^1$. The actor has quadratic loss preferences around her ideal outcome (o_i^*). A policy mapping function $\psi(p) \in \Psi$ maps each policy (p) to an outcome (o) (see below). This policy mapping function, and thus the policy uncertainty, is a Brownian motion. Our decision maker knows only that the policy map is a Brownian motion with drift μ and variance σ^2.[1]

The actor learns one policy-outcome pair by observing another's policy and outcome. [2] We also assume that our decision maker knows the other's ideal outcome and the distance (the difference) between their respective ideals (Δ_{ij}).

The full model (Glick, 2010) incorporates multiple actors, information costs, and other variables. It gives actors four options. The two we focus on are "mimicking" in which one adopts another's policy and "modifying" where one changes the other's policy. We elaborate predictions for the "mimicking" and "modifying" conditions and directly test those predic-

[1] In this application Brownian motion does not capture movement through time. There is no time element but the mathematics of physical time processes instead create a "policy mapping."

[2] Learning from one's own experience can be subsumed in this model. In this case, the known policy-outcome pair would be one that is known to the actor. Nevertheless, we conceptualize this as a model of learning from others. Most of the problems we are thinking of are ones in which it is too expensive for an individual to try different actions. Instead, they are trying to learn from others who have gone before to make their best one-shot decision. Nevertheless, Callander's papers which focus on policy experimentation incorporates some of the same math and logic while focusing on learning from one's own experiences when opportunities for experimentation are scarce leading to "stuck" policies.

tions with individuals in the lab. We investigate, theoretically and experimentally, when an individual will simply copy another's action and when (and how) they will alter it. The specific implications we tested in the lab follow directly from the Brownian motion uncertainty representation. Because this formalization is both novel and critical, and because one of our contributions is testing humans' ability to work with Brownian signals, we elaborate on it before deriving the key behavioral predictions.

3.1 Uncertainty as Brownian Motion

We discussed the evolution of these uncertainty models above along with some of the intuition. Nevertheless, because the Brownian motion assumption is so critical to our predictions, and because a central contribution of our experiment is a test of it, we now further elaborate on the mathematics. Recall that the canonical policy process assumes a mapping in which a policy plus a linear shock ω (drawn from a uniform distribution $[-\lambda, \lambda]$) results in an outcome. Formally, a policy p produces an outcome $o = p + \omega$. Once ω is revealed, the policy to produce an outcome 0_1 is just $p = 0_1 - \omega$.

We assume a different policy process in which $\psi(p)$ is a realization of a Brownian motion with drift parameter μ and variance σ^2. We assume that actors know the underling linear trend and the variance around it. (See figure 3.1 above). Learning a policy-outcome pair pins down the map at one point. After observing one policy mapping, i.e. policy p_j produces outcome o_j ($\psi(p_j) = o_j$), an actor has updated and improved beliefs about other policies. Specifically, knowing that $\psi(p_j) = o_j$, another policy p's expected outcome and its variance are:

$$\text{Expected Outcome: } E[\psi(p)] = o_j + \mu(p - p_j) \tag{1}$$

$$\text{Variance: } var[\psi(p)] = |p - p_j|\sigma^2 \tag{2}$$

This model of policy uncertainty is not only mathematically tractable, but has an intuitive interpretation (Callander, 2008). The expected value of policy x is just the slope of the line (the drift) multiplied by the distance from the known point to policy p. The variance is growing proportionally to the distance to the unknown policy p. There is more uncertainty the further away one moves from a well understood policy to set a new one. The ratio $\frac{\sigma^2}{|\mu|}$

indicates policy complexity. The larger the ratio, the less one learns about the full mapping from observing one policy-outcome pair. In the limit, the process approaches either full invertibility or non-invertibility.

Figure 1: Example of Brownian motion of slope μ with one known policy-outcome pair (p_1, o_1)

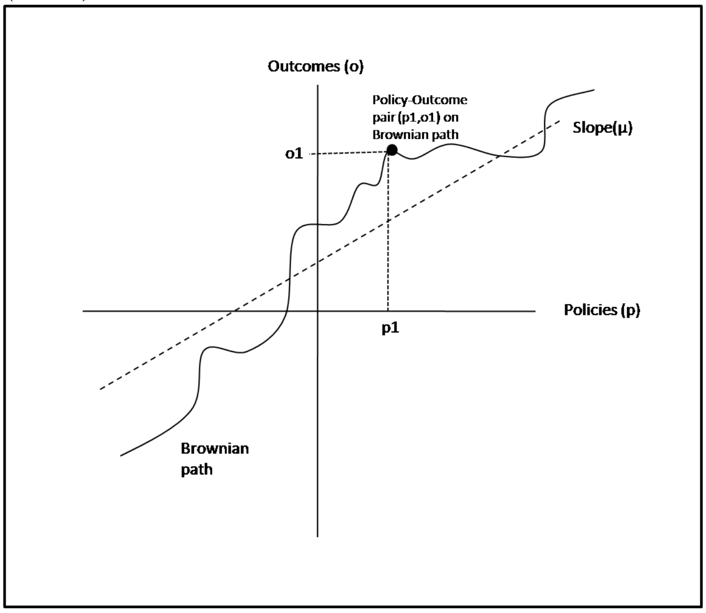

3.2 Mimicking vs. Modifying

We now briefly introduce mimicking and modifying solving for their respective expected utilities and policy outcomes in the two player setup as described above. Recall, we are interested in the later actor's behavior after learning from the firsts. This analysis parallels the analysis of "sticking" versus "experimenting" in Callander's (2011a) paper when beliefs are "open ended." We then analyze the tradeoff between the two ways of learning from others to derive the primary indifference condition we test. This analysis yields propositions concerning when one will mimic and when and how one will modify. We take a decision theoretic approach to these actions. This allows us to focus exclusively on this tradeoff which is a central mechanism in the model and which takes advantage of its novel uncertainty formalization. Moreover, as the full formal learning model shows, many of the possible incarnations, in both sequential and simultaneous choice settings, collapse into decision theoretic problems for the second player. In the future we would like to investigate some of the extensions that a game theoretic approach would support. These include strategic delay and free riding off of others' experiences, and commonality preferences in the future. These extensions significantly complicate matters. The cleaner decision theoretic setup approximates many real choices and we first examined the mimic vs. modify tradeoff in isolation.

1. **Mimic:** This option is straightforward copying. One implements another's policy (p_j) and gets their outcome. Since we are assuming that our actor observes the other's choice and outcome, this is now a known policy which produces a known outcome in an otherwise uncertain world. We assume (though it is not essential) that actor one actually identified their optimal policy and thus that the known policy produces the other's ideal outcome o_j^*.[3] Thus, mimicking produces an outcome exactly $|\Delta_{ij}|$ away from one's own ideal point because since it produces the others' ideal outcome and $|\Delta_{ij}|$ is the magnitude of the distance between their goals. If the two have exactly the same goals following is an easy choice.

$$\begin{aligned} EU_i(mimick_j) &= -\,\gamma_i \; E[(o_j^* - o_i^*)^2] \\ EU_i(mimick_j) &= -\,\gamma_i \Delta_{ij}^2 \end{aligned} \tag{3}$$

2. **Modify:** In this case, the actors observes the other's policy (p_j) and outcome (o_j) like before, but then chooses a different policy. Observing the first still provides information and reduces uncertainty and enables better "modified" policies. Here, the actor attempts to implement a policy closer to its own ideal point after incompletely learning about the policy map. (Recall, the expected outcome and utility of policies with a known policy-outcome pair are above.) We conceptualize this as starting with the others' policy and then making changes to it. For example, an instructor teaching

[3] This assumption is a strong simplification. Surely, there will be idiosyncratic variation around the known outcome. We would likely approximate this implementation variation as a draw from a mean-zero distribution (something like a uniform "ω shock"). Incorporating it would shift the cutoff towards modifying, but would not undermine the logic unless this variation was larger than the variance in the Brownian motion. More generally, the sharp cut point between modifying and mimicking is a simplifying assumption. In practice, both exist on a continuum where mimicking approximates slight modifying. These extensions, complications, and this type of variation are interesting avenues for future analysis, but the simple model captures the logic we'd like to explore as a first step.

an introductory class about congressional politics for the first time might find a more advanced syllabus and then replace some of the more technical works with simpler ones. This tactic and analysis are very similar to "open ended beliefs" in Callander's experimentation model.

$$
\begin{aligned}
EU_i(modify_j) =& -\gamma_i E[(o_i - o_i^*)^2] \\
EU_i(modify_j) =& -\gamma_i [E[o_i] - o_i^*)^2 + Var(o_i)] \\
EU_i(modify_j) =& -\gamma_i [(o_j^* + \mu(p_i - p_j) - o_i^*)^2 + |p_i - p_j|\sigma^2] \\
EU_i(modify_j) =& -\gamma_i [(\mu(p_i - p_j) - \Delta_{12})^2 + |p_i - p_j|\sigma^2]
\end{aligned}
$$

Without loss of generality, assume that μ and $(o_i^* - o_j^*)$ (the distance Δ_{ij} between i and j's ideal outcomes) are positive. Thus, the expected utility of implementing policy p_i which produces outcome $o_i = \psi(p_i)$ after observing that $\psi(p_j) = o_j^*$ is:

$$EU(p_i) = -\gamma_i[\mu(p_i - p_j) - \Delta_{ij}]^2 - \gamma_i(p_i - p_j)\sigma^2 \quad (4)$$

This is the general expression for the expected utility of altering. When altering, an actor still has choice about which policy to implement. Thus, we must solve for the expected utility of the optimal altered policy. This is the utility which in theory it should compare to the utility of mimicking. Deriving it will allow us to test both the when, and how well, questions about modifying. We denote the optimal altered policy $\dot{p}_i^*$ and find the p_i that maximizes equation 4. The first and second derivatives with respect to p_i are:

$$
\begin{aligned}
\frac{dEU}{dp_i} &= 2\mu\gamma_i[\Delta_{ij} - \mu(p_i - p_j)] - \gamma_i\sigma^2 \\
\frac{d^2EU}{dp_i^2} &= -2\gamma_i\mu^2
\end{aligned}
$$

Solving for p_i to get $\dot{p}_i^*$, the best altered policy given available information:

$$
\begin{aligned}
2\mu^2\dot{p}_i^* &= 2\mu^2 p_j^* + 2\mu^2\Delta_{ij} - \sigma^2 \\
\dot{p}_i^* &= p_j + \frac{\Delta_{ij}}{\mu} - \frac{\sigma^2}{2\mu^2}
\end{aligned}
\tag{5}
$$

This analysis leads to a precise prediction about the second actors optimal modified policy and this predictions sheds some light on the mechanics of decision making in this environment,

Proposition I: When modifying, choose optimal policy according to equation 5. Move in the direction of ideal policy but by an amount less than $|\Delta|$ (the difference in goals). This fraction of $|\Delta|$ will decrease with issue complexity.

Substantively, the optimal modified policy will be closer to the well understood one than it would be without uncertainty. The $p_j + \frac{\Delta_{ij}}{\mu}$ component is exactly what one would do to get to o_i^* if the policy mapping was linear with slope μ. Subtracting the variance term moves the optimal altered policy closer to the known one. Since the variance is multiplied by the distance between the policies, to reduce uncertainty (and expected loss), the gap should be smaller than linear extrapolation would imply. The variance term is proportional to issue complexity. As it increases, that is, when issues are more complex, the less one should modify an existing policy. In sum, altered policies should be relatively conservative (closer to the known policy than the ideal point gap suggests they would be without uncertainty) particularly when issues are more complex.

Next, we can then solve for the expected utility implementing this optimal modified policy $\dot{p}_i^*$. We substitute $\dot{p}_i^*$ from equation 5 into equation 4 to get:

$$
\begin{aligned}
EU(\dot{p}_i^*) &= -\gamma_i[\mu(p_j + \frac{\Delta_{ij}}{\mu} + \frac{\sigma^2}{2\mu^2} - p_j) - \Delta_{ij}]^2 - \gamma_i[p_j + \frac{\Delta_{ij}}{\mu} - \frac{\sigma^2}{2\mu^2} - p_j]\sigma^2 \\
EU(\dot{p}_i^*) &= -\gamma_i[\frac{\sigma^2}{2\mu}]^2 - \gamma_i[\frac{\Delta_{ij}}{\mu} - \frac{\sigma^2}{2\mu^2}]\sigma^2
\end{aligned}
$$

Thus, the expected utility of implementing the best modified policy is:

$$EU(\dot{p_i^*}) = -\gamma_i[\frac{\Delta_{ij}\sigma^2}{\mu} - \frac{\sigma^4}{4\mu^2}] \tag{6}$$

Having expressions for the expected utility of the two options, we can solve for the key indifference condition between them. This condition is the theoretical answer to the question **when is modifying the known policy better than mimicking it?** Formally, we do not need the "mimic" option and could just consider cases when a profitable modification does or does not exist or we could just think about a spectrum ranging from no modifications to large ones. Nevertheless, it may be more intuitive to at least think about choosing between taking something straight off the shelf or significantly modifying it rather than realizing that the optimal modification in a given case is zero (or very small). The easiest way to solve for the indifference is to revisit the optimal p_1^* equation (equation 5):[4]

$$p_2^* = p_1 + \frac{\Delta_{21}}{\mu} - \frac{\sigma^2}{2\mu^2}$$

Because μ and Δ_{21} are positive by construction, p_2 must be greater than p_1. This can only be satisfied, implying a gain to modifying, when $\frac{\Delta_{21}}{\mu} \geq \frac{\sigma^2}{2\mu^2}$. Thus, **modifying is preferred to mimicking when:**

$$\Delta_{21} \geq \frac{\sigma^2}{2\mu} \tag{7}$$

When complexity is large enough, or Δ is small enough, altering is no better than following. (The best altered policy will be zero alteration producing the "mimic" result). Modifying, or perhaps more precisely, modifying more rather than less, is more attractive relative to following (or modifying only a little) when goals are very different (large distance between the two actors' ideals), and when issues are less complex (small $|\frac{\sigma^2}{\mu}|$). The intuition

[4]We can see this (and reach the same indifference condition) by comparing equation 3 to equation 6. Rearranging we can see that the utility of the best modified policy is equal to the expected utility of the mimic option when $(\Delta_{ij} - \frac{\sigma^2}{2\mu})^2 = 0$. Thus, when $\Delta_{ij} = \frac{\sigma^2}{2\mu}$, the best modified policy is one which is not modified at all. Substituting this into the best modified policy equation just leaves p_j, the known policy, as the best "modified" policy.

is straightforward. A safe choice looks less and less attractive when it was made to satisfy very different goals (large Δ_{21}). On the other hand, at times, moving off a known policy, even in the expected correct direction, invites too much uncertainty. Sometimes, the potential for customization and improved fit is not worth the risk.

Proposition II: The likelihood of modifying a known policy, and the magnitude of modification will vary with the distance to the known policy and vary inversely with task complexity (7).

By means of contrast with some traditional models, compare this outcome to a model where the policy mapping function $\psi(p) \in \Psi$ is a simple linear shock such that $o_i = p_i + \mu$. In such a case, the optimal behavior is to always modify the policy, and adopt policy $p_j + \Delta_{21}$. Note that the difference between the two is not just that players always alter with a simple linear shock; the degree of alteration is always greater with the simple linear shock than with the more complex mapping function that renders the signal partially invertible.

4 Experimental Design

We tested the learning model above by giving participants partially invertible signals about a decision task. The model produces clear but somewhat counter-intuitive findings. Specifically, it suggests that in some cases actors' best action is to discard information about the difference between themselves and an informed actor by simply copying that actor's action. Even when actors make use of this information by modifying the other's policy, they should discount it and select a policy closer to the other's choice than a simple linear extrapolation would imply. Given the counter-intuitive nature of this theoretical finding, it is far from clear that actors will learn to throw seemingly useful information away or discount it to the optimal degree.

Further, determining the optimal action in this model depends on two different factors: the complexity of the policy decision and the similarity, or lack thereof, between policymakers. This is in contrast to the canonical model, where the second policymaker need only know the degree of similarity between his and the first actor's goals to act optimally. While both of these factors play a role in determining whether and by how much the second player should modify the first player's action, people may react differently to each source of uncertainty. This is particularly true given the fact that people in general do a poor job of understanding and estimating statistical variance and incorporating it into their judgments (Beach and Phillips, 1967; Garthwaite et al., 2005; Hogarth, 1975).

We presented participants with a simple decision task. In each round they had to pick a number. That number (policy / input) was subject to a stochastic shock to produce a new number (policy outcome / output). The participant's goal was to end up with a final number as close as possible to zero. Participants were informed that a randomly drawn number would be added to the number they chose to produce their new number. In each round we showed the participant a computer players' choice, the random number that was added to it, and the computer players' final outcome. Participants were then asked to choose a number, a random shock was drawn, and their final number was displayed. A new round then began and the participant was presented with another computer player's number, shock, and final number. Each participant completed 100 rounds. We randomly selected one round on which to base their earnings which depended on the squared distance of their final number (output) from zero.

If a participant choose the number that the computer player chose ("mimicking") they received the same outcome as the computer player. They knew that if they chose a different number, it would be subject to a different shock defined by the Brownian motion process.[5] We explained the stochastic component by explaining that the variability of the random number would grow the further they moved from the computer players' number. We showed them two sets of pictures for each variance (20 and 40) treatment to illustrate (figure 2). First, we showed them four histograms approximating the normal distributions from which the shock was drawn. These histograms corresponded to picking a number 5, 10, 15, and 20 units away from the computer player's. We also showed them a scatter plot with thousands of points drawn from the distribution. The distance from the computer player was on the x-axis and the shock was on the y. Subjects were given handouts with these figures, and could consult them during the experiment.

[5] Adding a small additional shock to both decisions would ensure that both decisions are subject to some uncertainty and add some realism, but we feel that doing so would complicate the experimental test without substantially altering the model being tested as it would just shift the cut indifference point.

Figure 2: Examples of information provided to participants to explain the stochastic process to them in the case where the variance was equal to 40

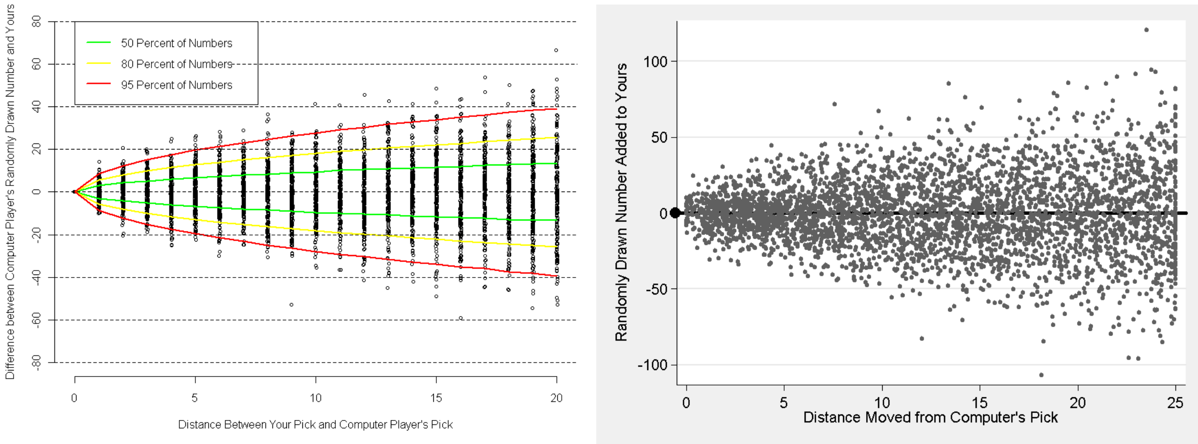

Deviations from equilibrium predictions may be a result of a failure to properly learn from the model. They may also be a result of risk aversion. Since choosing a policy closer to that chosen by the computer player will always produce an outcome with a lower expected variance, risk averse subjects will fail to modify to the degree predicted by the model even if they learn properly from the game. To account for this, we measured subjects' risk aversion using the gamble-selection task from (Eckel and Grossman, 2008). Subjects chose one of five gambles. Each gamble has two outcomes, with a fifty percent chance of each outcome. The first gamble was a sure thing - both resulted in the subject receiving \$4 - while the other four gambles increased both the expected payoff and the risk associated with that payoff in a linear manner. To make sure that a subject's measure of risk aversion was not affected by how that subject fared in the main experiment, Subjects selected a gamble prior to the beginning of the main experiment. The task was presented as a separate experiment from the main study; to further insulate the gamble selection from the main experiment instructions for the main study were given between the gamble selection and the main experiment. Their gamble selection was resolved after the experiment was completed when subjects were paid.[6]

We completed six sessions of the experiment with 67 subjects. One subject's data was dropped for failing to comply with experimental instructions leaving 20 subjects in condition

[6]Specifically, a coin was flipped to determine which payoff the subject would receive.

one, 12 in condition two, 16 in condition three and 19 in condition four. Subjects played the game for 100 rounds. We selected one round at random for payment. Subjects were paid \$10 minus the squared distance between their outcome in the randomly selected round and 0 in cents. They also received a \$10 show-up fee. Average earnings were \$26.81, including average earnings of \$5.23 from the gamble selection.

4.1 Experimental Predictions

In the model, two factors influence a player's decision to mimic or modify: the complexity of the decision task and the degree of similarity between the player's goals and the goals of the informed player. These are represented mathematically as the variance of the Brownian motion policy map and the distance between the players' ideal outcomes (Δ_{12}). To test participant learning from partially invertible signals we varied these two factors as experimental treatments. The participant's goal (optimal policy outcome) was fixed at zero. The distance between the participant's goal and the computer player's goal was either small ($\Delta_{12} = 10$) or large ($\Delta_{12} = 20$). The variance of Brownian motion process was also manipulated to make the task either complex ($\sigma^2 = 40$) or simple ($\sigma^2 = 20$). Table 1 shows these four experimental conditions. So that participants did not face seemingly identical tasks each time, the distance between the participant and the computer player was drawn from a uniform distribution on $(E(\Delta_{12}) - 5, E(\Delta_{12}) + 5)$. In all cases, the underlying trend or drift (μ) of the Brownian motion was 1.

		Complexity	
		High ($\sigma^2 = 40$)	Low ($\sigma^2 = 20$)
Similarity	High ($E(\Delta_{12}) = 10$)	(1)	(2)
	Low ($E(\Delta_{12}) = 20$)	(3)	(4)

Table 1: Experimental Conditions

The parameters were chosen to test the model's predictions that players should mimic when tasks are more complex or when similarity is high and modify when tasks are simpler or when similarity is low. The optimal policy choices for the experimental treatments are summarized in table 2.In condition (1), players should always mimic. In condition (4), players should always modify. In conditions (2) and (3) the expected value of Δ_{12} is on the cutpoint between altering and following. Depending on whether the draw of Δ_{12} is above or below the mean, players should either follow or alter. Thus, optimal behavior would lead them to mimic half of the time, and modify half of the time; the average modification would be five units.

Comparing actual behavior to these predictions will show how subjects respond to different kinds of difficulty in learning from others. By comparing condition (1) to condition (2) and condition (3) to condition (4) we can determine how the complexity of a task affects optimal decision making. By comparing condition (1) to condition (3) and condition (2) to condition (4) we can determine how the similarity (or lack thereof) between decision makers affects optimal decision making. Finally, comparisons between subjects in conditions (2) and (3), and how behavior in these conditions differs from behavior in conditions (1) and (4) will allow us to tell whether participants respond in different ways to these two types of task difficulty.

		Complexity	
		High ($\sigma^2 = 40$)	**Low** ($\sigma^2 = 20$)
Similarity	**High** ($E(\Delta_{12}) = 10$)	p_j	p_j if $\Delta_{12} < 10$ $p_j - \Delta_{12} + 10$ if $\Delta_{12} > 10$
	Low ($E(\Delta_{12}) = 20$)	p_j if $\Delta_{12} < 20$ $p_j - \Delta_{12} + 20$ if $\Delta_{12} > 20$	$p_j - \Delta_{12} + 10$

Table 2: Equilibrium Predictions

5 Results

We focus on two metrics to measure how subject behavior compares to theoretical predictions. The first is deviation from the optimal action - how many units the player's actions falls from the optional action, as defined in Table 2. A positive deviation indicates over-modifying the first player's action, while a negative deviation indicates that the subject was too cautious, under-adjusting the first player's action compared to optimal behavior. The second metric is how frequently the subject mimicked the actions of the computer player. According to Table 2, subjects should always mimic in condition 1, should mimic fifty percent of the time in conditions 2 and 3, and should never mimic in condition 4. Translated into comparative static predictions, subjects should mimic the most in condition 1, mimic less frequently in conditions 2 and 3 than in condition 1, and mimic the least in condition 4.

In general, our results suggest two things: a bias towards action vs. inaction, and a tendency to over-modify in high-complexity situations. Figure 3 shows the average deviation from the optimal modification for each condition over time, grouped into ten-period blocks. This figure shows that subjects performed worst in conditions where the complexity of the task was high (1 and 3). In both of high-complexity conditions the mean decision modified the computer player's action by considerably more than was optimal, whereas in conditions 2 and 4 subjects appear to have learned to make, on average, the correct modification. Looking at the effect of similarity yields more mixed results. Subjects in high-similarity condition 1 did better than those in the comparable low-similarity condition 3. However, subjects in low-similarity condition 4 did slightly better than subjects in low-similarity condition 2.

Figure 4 shows the overtime trends in the percent of subjects mimicking in each condition. Recall that optimal behavior would have subjects mimic all of the time in condition 1, half the time in conditions 2 and 3, and never in condition 4. Subjects come closest to these behaviors in condition 4, where they almost never mimic at any point in the experiment. In the other three conditions subjects behave optimally a bit more than half the time, and do a bit worse in the low-similarity high complexity condition (3) than the high-similarity low-complexity condition (2). These results are roughly in line with the comparative static predictions described above. The offer some evidence suggesting that subject's came closer to optimal behavior in low-complexity conditions than in high-complexity conditions. Overall, the strongest suggestion is a bias towards action, as subjects did mimicked less than they should have in all three conditions where mimicking could be an optimal action.

Examining these results systematically requires moving from examining individual decisions to examining the behavior of subjects aggregated over time. Since subjects learn from round to round, the individual decisions taken by a single subject are likely to be highly correlated. Thus, any statistical test that used the individual decision as the unit of analysis would badly understate its standard errors; examining the behavior of individual subjects avoids for this issue. We characterize each subject by taking the percent of times that each

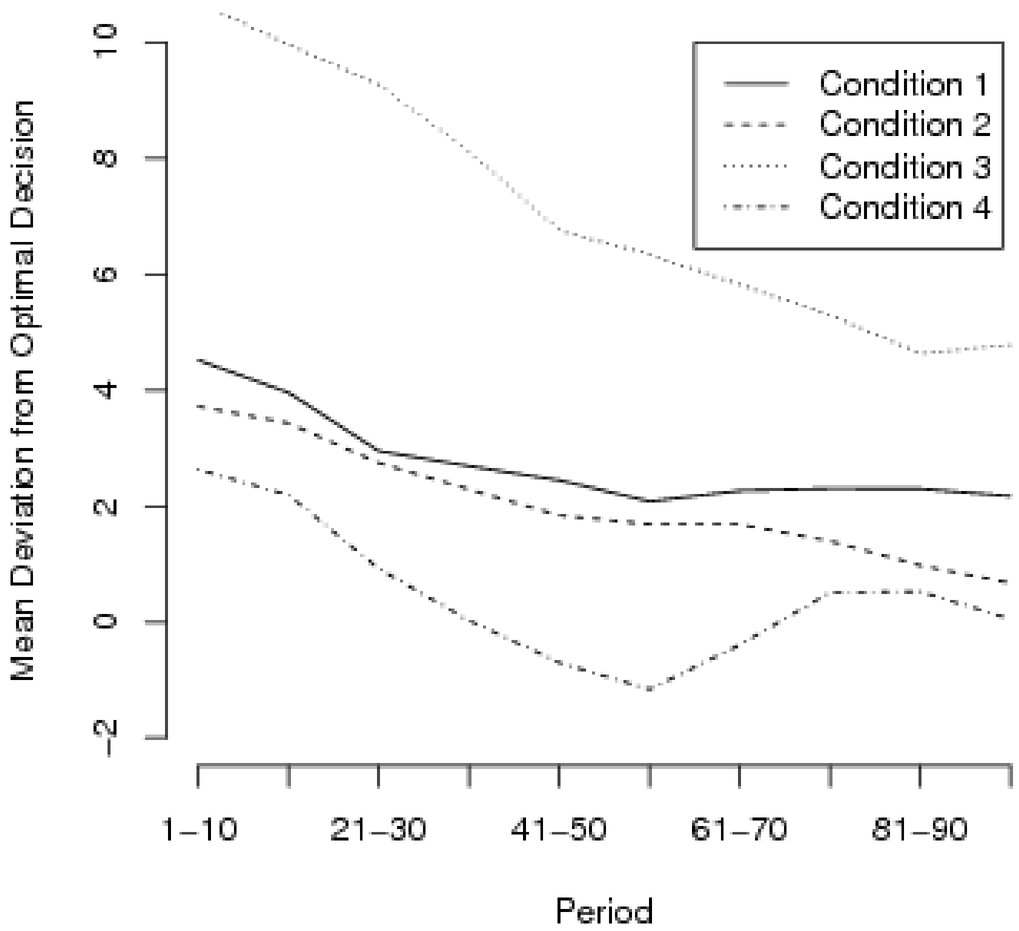

Figure 3: Mean Deviation from Optimal Action by 10 Period Block

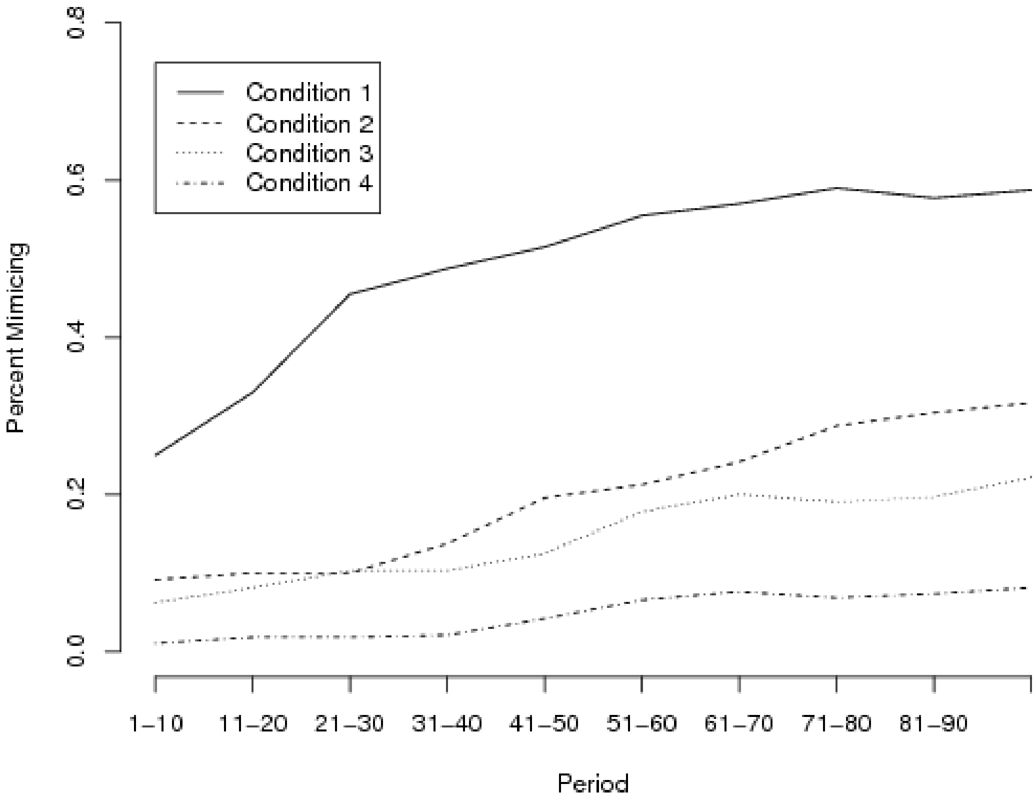

Figure 4: Percent of Mimicking Decisions by 10 Period Block

subject mimicked and the subject's average deviation from the optimal action across the final 20 rounds of play, a period when it appears that average behavior has stabilized.[7] We then compare the distributions of these variables across experimental conditions.

The histograms in Figure 5 tends to confirm the inferences drawn above. These figures show the distribution of subjects' average deviation from the optimal action, where a negative deviation signifies modifying less than is optimal while a positive deviation signifies modifying more than is optimal. This figure shows that that the subjects in low-complexity conditions tended towards optimal behavior while the average subject in high-complexity conditions tended to modify too much. Differences in similarity did not change the central tendency of subject behavior. However, there is also a great deal of heterogeneity across subjects, and lower similarity did increase this heterogeneity.

Examining the percentage of times that subjects mimicked again shows a bias towards action in all conditions where mimicking could be an optimal action. In the high-similarity, high-complexity condition, where the optimal action was to always mimic, A majority mimicked less than 80 percent of the time. In the two intermediate conditions (conditions 2 and 3) where the optimal action would call for mimicking roughly half the time almost no subjects behaved in this manner. Instead, subjects tended to mimic either almost always or almost never. Comparing these conditions, subjects in the low-similarity, high complexity condition tended towards less mimicking, again showing a failure to adjust properly to the effects of increased complexity.

To confirm these observations, Table 3 shows the results of statistical tests comparing behavior in each condition to each other condition, as well as to optimal behavior. The first four columns report two-sample t-tests comparing the average deviation in the condition listed in the row to the cognition listed in the column.[8] In general, the differences between conditions with the same level of similarity but different levels of complexity are significant, while the differences between conditions with the same level of complexity but different

[7] We have also examined the final 10 an 30 rounds of play, with substantively similar results.

[8] Two non-parametric tests, not reported here, yield substantively similar results. The first in a Wilcoxon Rank Sum test, a non-parametric analog of the t-test. The second is a two sample Kolmogorov-Smirnov test, which is conceptually different and tests whether the two samples were drawn from the same distribution. The latter test is important, as it examines whether the distribution of subject behavior is different in different conditions, not just the mean subject behavior.

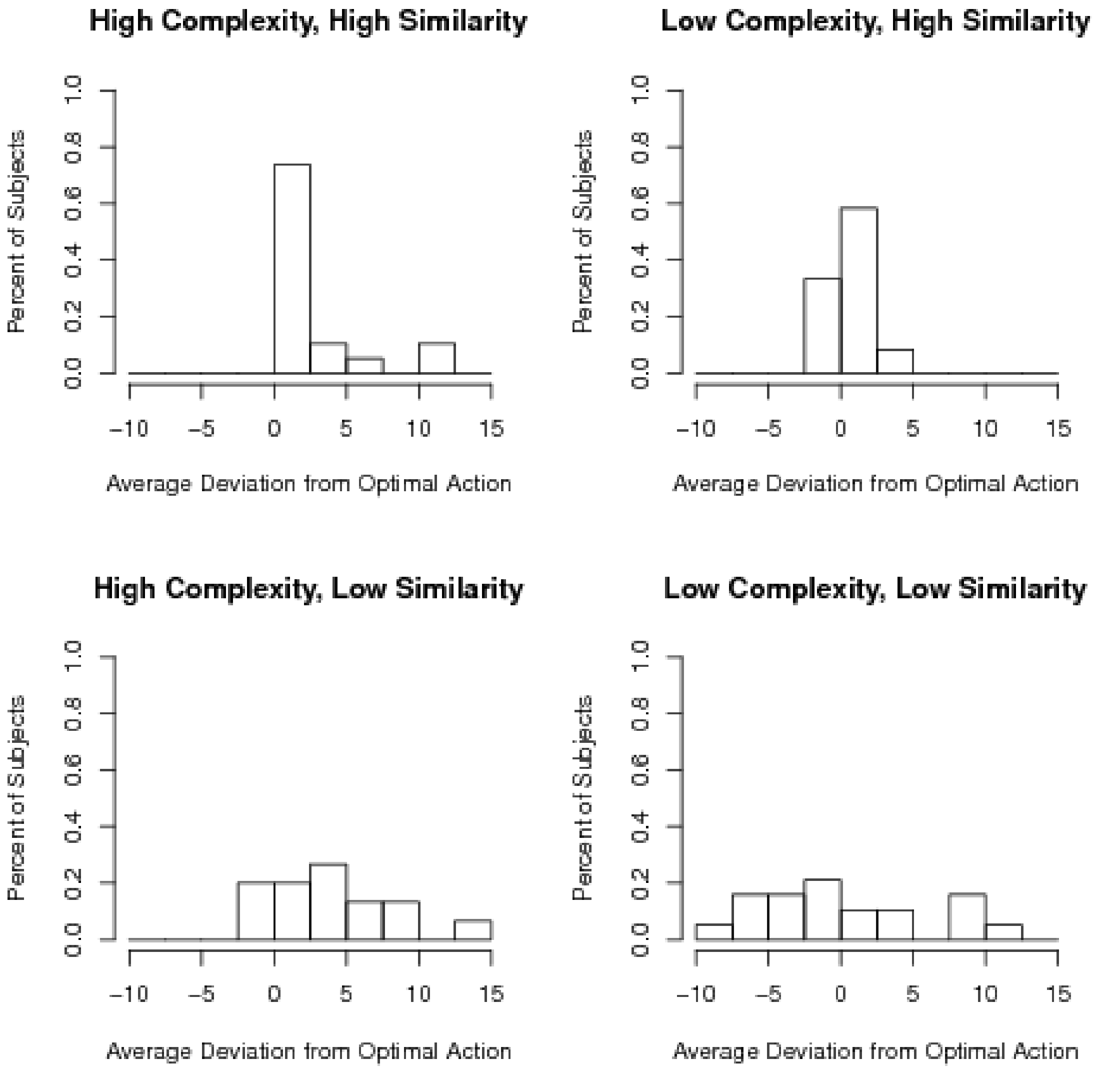

Figure 5: Subjects' Average Deviation from Optimal Action, Last 20 Rounds

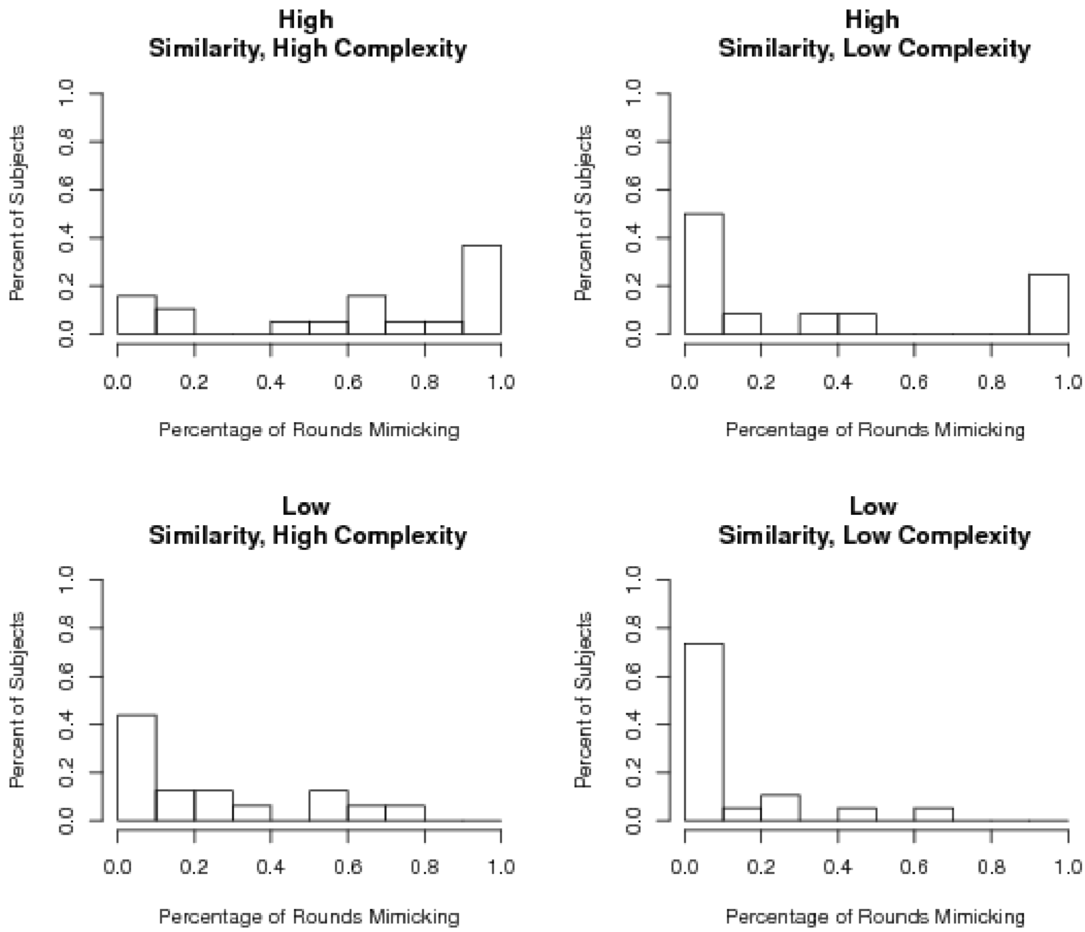

Figure 6: Subjects' Percent of Decisions that Mimic, Last 20 Rounds

	Low-Complexity High-Similarity	High-Complexity Low-Similarity	Low-Complexity Low-Similarity	Optimal Action	Optimal Mimicking
High-Complexity High-Similarity	$\Delta = 1.51$ ($p = .081$)	$\Delta = 2.61$ ($p = .100$)	$\Delta = 2.12$ ($p = .197$)	$\Delta = 2.17$ ($p = .015$)	$\Delta = -0.38$ ($p = 0.000$)
Low-Complexity High-Similarity	X X	$\Delta = 4.11$ ($p = .001$)	$\Delta = 0.62$ ($p = .674$)	$\Delta = 0.67$ ($p = .170$)	$\Delta = -0.23$ ($p = 0.162$)
High-Complexity Low-Similarity	X X	X X	$\Delta = 4.72$ ($p = .018$)	$\Delta = 4.78$ ($p = .002$)	$\Delta = -0.28$ ($p = .001$)
Low-Complexity Low-Similarity	X X	X X	X X	$\Delta = 0.05$ ($p = .971$)	$\Delta = 0.08$ ($p = .056$)

Table 3: Statistical Tests of Differences, Showing Differences and T-Test p Values

levels of similarity are not significant. Looking at the comparison with optimal behavior, the two high-complexity conditions have average deviations that are statistically significant from 0, while this difference is much smaller and not significant for the low-complexity conditions. Both high complexity conditions have a statistically significant deviation from the optimal rate of mimicking; among the low-complexity conditions only the low-similarity condition has a significant deviation from the optimal rate of mimicking, and this deviation is substantively fairly small. High complexity raises the average deviation from optimal behavior, while low similarity has no comparable effect.

6 Discussion and Conclusion

Callander's model offers a more nuanced, flexible, and theoretically attractive representation of uncertainty. It has applications to institutions and individuals in political settings and beyond. However, behaving consistent with equilibrium predictions predicated on this underlying uncertainty model requires more cognitive capacity from decision makers than previous models. In the canonical model, actors are faced with a fairly simple task. In contrast, modeling uncertainty as Browning motion presents actors with two challenges. The challenge posed by the similarity (or lack thereof) between a known policy and an unknown policy, and the challenge posed by the complexity of the policy domain. As the theoretical discussion in Section 3 shows, these challenges combine to make optimal action difficult. Further, they produce the counter-intuitive finding that sometimes actors are best off discarding information and simply mimicking the known policy. Given these complexities and the models' potential, it is important to know how actors actually behave when facing these two different challenges and whether they will mimic when this counterintuitive action is optimal.

Our data suggests that subjects fail to fully take into account the effect of high task-complexity. In low complexity conditions, aggregate data shows subjects doing a fairly good job of approximating the optimal action. However, while the average subject acts optimally, there is a great deal of variation with some subjects under-modifying and others over-modifying. In high-complexity conditions subjects show a bias towards over-modifying, with the average subject changing the observed policy considerably more than is optimal. This reflects a tendency to underestimate the effect that complexity in limiting the upside of large modifications policies. In other words, in high complexity environments, subjects invited extra risk. Finally, subjects show a bias towards action. They mimic only somewhat more than half as often as they should. This tendency was stronger in high-complexity conditions. All of these tendencies appear even after considerable opportunity to learn, and appear to be stable towards the final rounds of the game. In sum, subjects fail to take into account the complexity of the decision into account, and as a result modify too much. They are only partly able to incorporate partially invertible signals in the sophisticated ways that theory based on the new uncertainty representation predicts. Some actors may not be able to meet the cognitive demands of the model, and behave differently than its equilibrium predictions. These findings suggest some caution in applying the model as well as future work to examine when it is most applicable and what can be done to enhance individuals ability to learn optimally.

References

T.K. Ahn, Robert Huckfeldt, and Alex Mayer. Political experts, communication dominance, and patterns of political bias. *Midwest Political Science Association Annual Meeting*, 2009.

Brady Baybeck, William D. Berry, and David A. Siegel. A strategic theory of policy diffusion via intergovernmental competition. *The Journal of Politics*, 73(1):232–247, 2011.

Lee R. Beach and Lawrence D. Phillips. Subjective probabilities inferred from estimates and bets. *Journal of the American Experimental Psychology*, 75(3):354–359, 1967.

Bernard Berelson, Paul Lazarsfeld, and William McPhee. *Voting: A Study of Public Opinion Formation in a Presidential Campaign*. University of Chicago Press, Chicago, 1954.

S. Bikhchandani, D. Hirshleifer, and I. Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5):992, 1992.

S. Bikhchandani, D. Hirshleifer, and I. Welch. Learning from the behavior of others: Conformity, fads, and informational cascades. *Journal of Economic Perspectives*, 12:151–170, 1998.

Steven Callander. A theory of policy expertise. *Quarterly Journal of Political Science*, 3(2), 2008.

Steven Callander. Searching and learning by trial-and-error. *American Economic Review*, 101, 2011a.

Steven Callander. Searching for good policies. *American Political Science Review*, 105(2), 2011b.

Frank Dobbin, Beth Simmons, and Geoffrey Garrett. The global diffusion of public policies: Social construction, coercion, competition, or learning. *Annual Review of Sociology*, 33: 449–472, 2007.

Anthony Downs. *An Economic Theory of Democracy*. Harper, New York, 1957.

Catherine C. Eckel and Philip J. Grossman. Forecasting risk attitudes: An experimental study using actual and forecast gamble choices. *Journal of Economic Behavior and Organization*, 68(1):1–17, 2008.

Zachary Elkins and Beth Simmons. On waves, clusters, and diffusion: A conceptual framework. *The Annals of the American Academy of Political and Social Science*, 598(1):33, 2005.

Paul H Garthwaite, Joseph B Kadane, and Anthony O'Hagan. Statistical methods for eliciting probability distributions. *Journal of the American Statistical Association*, 100 (470):680–701, 2005.

Thomas W. Gilligan and Keith Krehbiel. Collective decisionmaking and standing committees: An informational rationale for restrictive amendment procedures. *Journal of Law, Economics, Organization*, 3(2):287–335, 1987.

Thomas W. Gilligan and Keith Krehbiel. Asymmetric information and legislative rules with a heterogeneous. *American Journal of Political Science*, 33(2):459–490, 1989.

David Glick. Mimicking, modifying, and policy diffusion: A formal model of learning with interview evidence from legal implementation. *Working Paper*, 2010.

Lawrence J. Grossback, Sean Nicholson-Crotty, and David A. M. Peterson. Ideology and learning in policy diffusion. *American Politics Research*, 32(5), 2004.

Robin M. Hogarth. Cognitive processes and the assessment of subjective probability distributions. *Journal of the American Statistical Association*, 70(350):271–289, 1975.

Robert Huckfeldt and John D. Sprague. *Citizens, Politics, and Social Communication: Information and influence in an election campaign.* Cambridge University Press, Cambridge, 1995.

Andrew Karch. Emerging issues and future directions in state policy diffusion research. *State Politics and Policy Quarterly*, 7(1), 2007.

Arthur Lupia and Matthew McCubbins. *The Democratic Dilemma: Can Citizens Learn What They Need to Know?* Cambridge University Press, Cambridge, 1998.

Cavadonga Meseguer. What role for learning? the diffusion of privatisation in oecd and latin american countries. *Journal of Public Policy*, 24(03):299–325, 2004.

Cavadonga Meseguer. Rational learning and bounded learning in the diffusion of policy innovations. *Rationality and Society*, 18(1):35, 2006.

Charles R. Shipan and Craig Volden. The mechanisms of policy diffusion. *American Journal of Political Science*, 52(4):840–857, 2008.

David Strang and Mary Still. In search of the elite: revising a model of adaptive emulation with evidence from benchmarking team. *Industrial and Corporate Change*, 13(2):309–333, 2004.

Jean-Robert Tyran and Rupert Sausgruber. The diffusion of innovations - an experimental investigation. *Journal of Evolutionary Economics*, 15(4):423–442, 2005.

Craig Volden. States as policy laboratories: Emulating success in the children's health insurance program. *American Journal of Political Science*, 50(2):294–312, 2006.

Craig Volden, Michael M. Ting, and Daniel P. Carpenter. A formal model of learning and policy diffusion. *American Political Science Review*, 102(3):319–332, 2008.

Kurt Weyland. Theories of policy diffusion: Lessons from latin american pension reform. *World Politics*, 57(2):262–295, 2005.