

What Makes a Rule Complex?

By RYAN OPREA*

We study the complexity of rules by paying experimental subjects to implement a series of algorithms and then eliciting their willingness-to-pay to avoid implementing them again in the future. The design allows us to examine hypotheses from the theoretical “automata” literature about the characteristics of rules that generate complexity costs. We find substantial aversion to complexity and a number of regularities in the characteristics of rules that make them complex and costly for subjects. Experience with a rule, the way a rule is represented, and the context in which a rule is implemented (mentally versus physically) also influence complexity.

When people fail to behave as economists expect them to, they tend to err on the side of simplicity.¹ A long tradition in economics (e.g. Simon, 1955) argues that this is, in part, driven by the fact that people dislike (or are incapable of) implementing complex rules and procedures. In this paper we propose new experimental methods to identify what makes a rule complex to implement, and to measure the costs of this “procedural complexity.” Doing this may allow us to better understand why people don’t always choose to use optimal procedures (e.g. strategies, dynamic optimization policies, state-contingent agreements, belief updating rules, choice procedures) or perfectly comply with rules handed down by governments and organizations (e.g. regulations, laws, bureaucratic rules).

Our investigation is built around the idea that rules are in fact *algorithms*,

* Economics Department, University of California, Santa Barbara, Santa Barbara CA, 93106-9210, *Email: roprea@gmail.com.* I thank Youssef Benzarti, Peter Boessaerts, Andy Brownback, Mark Dean, Akshay Dixit, Guillaume Frechette, David Freeman, Tamsin German, Daniel Lokshtanov, Kirby Nielsen, John Rehbeck, Collin Raymond, Julian Romero, Yaroslav Rosokha, Hamid Sabourian, Adam Sanjurjo, Charles Sprenger, Emanuel Vespa and Sevgi Yuksel for their detailed comments and discussions. I also thank Lucas Coffman, Ben Enke, Yoram Halevy, PJ Healy, Daniel Martin and Dale Stahl for valuable conversations about this research. I am grateful to seminar audiences at the University of Arkansas, the University of California, Davis, the University of Melbourne, Purdue University, the University of California, Santa Barbara, the University of Texas, Austin, the University of California, Berkeley, the University of California, San Diego and audiences at the Stanford Institute for Theoretical Economics, the 2019 Society for the Advancement of Economic Theory meetings in Ischia, the AFTS Workshop at University of Adelaide and the 2019 Economic Science Association North American Meetings in Los Angeles for their comments. Finally, I thank the co-editor, Stefano DellaVigna, and four referees for their constructive comments and guidance. This research was approved by UC Santa Barbara IRB and was supported by the National Science Foundation under Grant SES-1949366.

¹ For instance, many key tendencies from behavioral economics seem to involve simplification: satiety simplifies rational choice by radically reducing the number of rank comparisons made; myopia simplifies dynamic optimization by ignoring the impact of current acts on future outcomes; representativeness simplifies Bayes’ rule by ignoring the influence of base rates on posteriors; narrow bracketing simplifies joint optimization by ignoring the relationships between choices; adaptive learning simplifies rational expectations by ignoring counterfactuals; the winner’s curse and related errors simplify Bayesian Nash Equilibrium by ignoring contingencies or information in others’ choices etc.

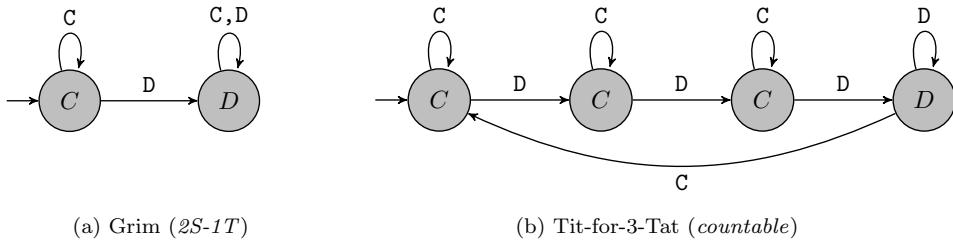


Figure 1. : Automata representations of two rules (strategies in the repeated prisoner's dilemma).

Note: In parentheses are names for the equivalent automata used in the experimental design.

implemented not by computers but by human actors and thinkers. The “automata literature” (surveyed in Kalai (1990); Chatterjee and Sabourian (2009)) formalizes this idea by mathematically describing rules as “automata,” formal descriptions of algorithms from computer science. This literature advances several hypotheses (reviewed in Section I) about what structural characteristics of automata make the rules they describe *complex*, operationalizing “complexity” as the *cost* of implementing a rule. We use these hypotheses – and the literature’s operationalization of complexity as cost – to organize our investigation.

Figure 1 shows examples of automata from the repeated games literature, the first place this idea was applied in economics. Panel (a) shows the automaton description of a rule called *Grim* (the “grim trigger strategy”), a *strategy* for playing the repeated prisoner’s dilemma. The automaton consists of two *states* (shown as circles) and a set of *transitions* (shown as arcs). Letters inside the circles tell the player what to do in each state (e.g. C for “choose to cooperate”, D for “choose to defect”) and letters next to each arc tell the player when to follow the corresponding transition (e.g. when your counterpart chooses C, D). The rule instructs the player to begin by choosing C (cooperate), but then to transition to the second state and choose D (defect) as soon as the other player chooses D. The right state is *absorbing*, meaning no transitions lead away from it, so once here the player chooses D forever. Panel (b) of Figure 1 shows the automaton description of another strategy discussed in the literature called “Tit-for-3-tat,” (Tf3T), a lenient version of the famous tit-for-tat strategy. This rule instructs the player to cooperate until her counterpart has defected 3 times and then to switch to defect, forgiving her afterwards by starting the rule over (shown as an arc transitioning back to the first state) if she cooperates.²

²For expositional ease, these automata include only transitions that depend on other players’ actions – to completely describe a strategy, transitions must also depend on the player’s own actions (see Kalai and Stanford, 1988). Likewise, normally Tf3T requires the player to wait for three contiguous rounds of defection before defecting; for simplicity we show a cumulative version that doesn’t require the three rounds of defections to be contiguous.

Intuitively Tf3T seems more complex to implement than Grim, but in what sense is this true? What is it about a rule like Tf3T that makes it “more complex” than a rule like Grim, and how is this complexity formally reflected in its algorithmic structure?

To answer questions like this, we conduct an experiment (described in Section II) in which we assign subjects a sequence of abstract rules of the form e.g. “Choose x until you see a , after which switch to y ” and show them a sequence of randomly generated events (letters in the alphabet). The subject’s task is simply to correctly *implement* the rule in response to the random events by typing a sequence of letters on her keyboard. After subjects have implemented a number of different rules, we elicit their *willingness to pay* to avoid having to implement each rule again in the future, allowing us to measure the subjective cost of implementing each rule. By varying characteristics of the automata describing these rules, we can measure what algorithmic characteristics cause a rule to be complex and costly.

For instance, one popular theory from the automata literature is *state complexity* (or s -complexity), the hypothesis that states generate complexity costs. States describe the degree to which a rule-follower must condition her behavior on the past and they also summarize the distinct ways a rule requires a decision maker to process and respond to information. In our example, Tf3T contains two more states than Grim and by comparing how much subjects are willing to pay to avoid Tf3T versus Grim we can measure the subjective cost of adding two states to a rule, allowing us to evaluate s -complexity. In the experiment we find that subjects are willing to pay *twice as much* to avoid Tf3T as Grim, providing some evidence that states produce complexity costs. (Other natural complexity measures like implementation time are also generated by our data and we can perform the same exercise on them – for instance it takes subjects three times as long to implement Tf3T as Grim.)

Of course many algorithmic characteristics (not only states) differ between Tf3T and Grim and so the experiment is designed around a set of carefully chosen rules that independently vary automata characteristics to study their effects. For instance, Tf3T also contains many more *transitions* than Grim, and *transition complexity* (t -complexity, e.g. Banks and Sundaram (1990)) is the theory that this generates complexity in a rule. Transitions summarize the intensity with which a rule requires a decision maker to *monitor* and respond to the environment, summarizing a different aspect of a rule than states. Independently varying states and transitions and comparing behavior across our rules, we find that adding a transition indeed adds complexity costs, but that the effect is half as large as that of adding a state. Grim also has a simpler *architecture* than Tf3T because it contains an absorbing state and thus never requires the rule-follower to return to previously visited states. We find evidence that the architecture of rules shapes complexity too: *ceteris paribus*, rules with absorbing states are considerably less complex. Overall, our cost estimates suggest that rules like Tf3T are more com-

plex (costly) than rules like Grim and we therefore might expect them to be used less frequently. As we discuss in Section IV, this is in fact true in experimental repeated prisoner's dilemmas.

Our experiment also allows us to study how people *represent* rules to themselves and how this influences complexity. For instance, the literature generally assumes efficient representation: people don't over-complicate plans of action by representing them as rules containing unnecessary states. We test this by giving subjects similar rules that do and do not contain redundant states and we find strong evidence against efficient representation: subjects do not treat rules that contain redundant states as any less complex. A related but more fundamental assumption is that people only represent behavior using rules describable by *finite state machines* (FSMs), the basic class of algorithms studied by the automata literature. Tf3T, for example, cannot be represented as an FSM with fewer than four states, but if represented as a slightly richer type of algorithm (a "pushdown automata") that can simply *count* the number of 'D' events (by holding and manipulating a simple object in working memory), it can be represented with only two states and thus be dramatically simplified. When we compare responses to rules like Tf3T to rules that cannot be represented as counting rules, we find that indeed Tf3T is far less complex (less costly, faster to implement), suggesting that subjects are not constrained to represent behavioral patterns as FSM-like rules. Our results thus suggest that richer descriptions of algorithms than are typically studied in the automata literature are likely important for fully understanding and characterizing procedural complexity.

Finally, our experimental task is abstract but we take some first steps towards understanding how *context* influences the complexity of rules. For instance, we show that familiarity with a rule has a powerful alleviative effect on complexity: subjects given repeated exposure to a rule learn to implement it much faster and grow less willing to pay to avoid the rule. Likewise, a treatment variation in our design reveals that it is significantly more complex (from both a time and cost perspective) for subjects to enact rules mentally (e.g. to think through a counterfactual or forecast another's behavior) than it is to implement a procedure via a sequence of acts (e.g. to enact a strategy in a repeated game).

The purpose of our experiment is not to "test automata models" or even to propose automata as models of human cognition. Rather, automata are formal ways of describing and taxonomizing *rules*, and the purpose of our experiment is to study what (if any) components of these formal descriptions (e.g. states, transitions) systematically translate into complexity costs for decision makers, given the rules they describe. To the degree algorithmic characteristics predict complexity costs, we may be able to use automaton descriptions of rules, procedures, strategies etc. (combined with estimates using methods like ours) to model, predict and explain behavior in applications (e.g. to understand what sorts of rules people are unlikely to use or follow).

Overall, our findings (described in detail in Section III) can be summarized as

follows:

- 1) Most subjects find implementing complex rules costly, but these costs are highly heterogeneous across subjects and linked to independent measures of cognitive ability.
- 2) Algorithmic characteristics of rules significantly influence their complexity. Rules that require more **states** are significantly more complex (more costly, take longer to implement). Adding a **transition** is also costly, but adds half as much cost.
- 3) “Architectural” characteristics of algorithms other than simple state and transition counts influence complexity too: rules that eventually terminate in **absorbing** states (rather than endlessly doubling back to previously visited states) are significantly less complex (equivalent to removing a state from the rule).
- 4) Subjects have difficulty efficiently representing rules to themselves by discarding **redundant states**. However, subjects are able to represent some rules to themselves in efficient ways not achievable using the simplest types of algorithms usually studied in the literature. For instance, subjects seem able to conserve complexity by tracking simple **sequences in working memory**, allowing them to remove costly states from rules.
- 5) Complexity costs fall significantly as subjects become experienced at (familiar with) a rule, but this “**procedural learning**” is fragile and not easily transferred to even superficial perturbations of the learned rule.
- 6) The context in which a rule is implemented can matter a great deal for complexity. Employing a rule **mentally** (e.g. to make inferences) is significantly more costly than employing a rule physically (e.g. to implement a strategy).
- 7) The features of rules that generate cost are closely related to the features that generate implementation time, another important complexity metric. By contrast, there is only a very weak relationship between the drivers of costs and of implementation mistakes.

Although our motivating examples are strategies, we emphasize that *any* decision procedure or rule has algorithmic structure and so the methods and measurements we study here have potential applications across economics. Perhaps most intriguingly, aversion to procedural complexity has been used to explain (via automata models) important findings in *behavioral economics* including satisficing, primacy and recency effects, choice overload and status quo bias (Salant, 2011), stochastic choice (Kalai and Solan, 2003), non-Bayesian inference (Chauvin, 2020), biases in information processing (Wilson, 2014), and failures of back-

wards induction (Neyman, 1985).³ Understanding the degree to which distaste for complexity underlies and ties together classic findings from behavioral economics is a particularly important application for our methods. Earlier, automata were used to understand how aversion to complexity might solve equilibrium selection problems in games (e.g. Rubinstein, 1986; Abreu and Rubinstein, 1988; Kalai and Stanford, 1988), another promising application.⁴ A third potentially valuable application is to use methods like ours to study the algorithmic complexity of regulations, bureaucracies, laws, tax codes etc. in order to understand why policies and institutions don't always work as intended and to guide the design of more effective alternatives (see Abeler and Jäger (2015) for related work).

Though our focus is on understanding what makes *behavior* complex ("implementation complexity"), our methods may also be useful for understanding what makes some *decision problems* complex ("computational complexity"). This is because the *computational complexity* of a problem is typically operationalized (e.g. in computer science) as the *implementation complexity* of the most efficient evaluative procedure capable of solving it. In some cases, optimal decision making may be complex because it requires decision makers to use elaborate solution procedures that are complex to implement (costly, time-intensive, difficult) in a similar way and for similar reasons that other rules (such as those in our experiment) are.⁵ For example, Salant (2011) shows in an automata model that the computational problem of selecting an optimal option from a list is implementation-complex in exactly the sense we measure here. Several recent experimental papers (e.g. Carsten Murawski and Peter Bossaerts, 2016) have studied how the computational complexity of optimization problems impacts subjects' performance, and studying the degree to which implementation complexity measures like ours (perhaps elicited over more sophisticated types of algorithms like Turing machines) can be used to understand the effects of computational complexity on choice seems an important task for future research.

Finally, we operationalize complexity *behaviorally*, as a measure of the subjective burden a rule places on a decision maker. In doing this we follow several adjacent literatures. Computer science similarly defines complexity as a measure of the resources a machine must expend in order to implement an algorithm

³Recent experimental evidence suggests that subjects prefer to evaluate risk and lotteries procedurally (Halevy and Mayraz, 2020; Nielsen and Rehbeck, 2020) meaning procedural complexity may influence risky choice as well.

⁴Automata models of complexity have also been used in applied strategic contexts ranging from bargaining (e.g. Chatterjee and Sabourian, 2000) and industrial organization (e.g. Fershtman and Kalai, 1993; Sherstyuk, 2011) to general equilibrium (e.g. Sabourian, 2004; Gale and Sabourian, 2005).

⁵See Bossaerts and Murawski (2017) for a discussion of computational complexity and human behavior, emphasizing that even basic decision problems in economics can be intractably computationally complex. Echenique, Golovin and Wierman (2011) make the useful point that although optimization can be complex, simply satisfying rationality in the revealed preferences sense needn't be. Finally, it is worth emphasizing that some decision problems are "difficult" not because they require computationally intensive procedures but instead because the most attractive and immediate models of the problem or procedures for solving them are simply wrong. For instance, in the famous Monty Hall problem or the Wason selection task, the computations required to choose optimally are trivial but most people make fundamental mistakes because the intuitively appealing rule is faulty.

– usually time (“time complexity”) or working memory (“space complexity”). By focusing on the subjective costs of rules (“cost complexity”), we also follow growing literatures in both cognitive psychology (e.g. Wouter Kool and Matthew Botvinick, 2018; Amitai Shenhar, Sebastian Musslick, Falk Lieder, Wouter Kool, Thomas L. Griffiths, Jonathan D. Cohen and Matthew M. Botvinick, 2017). and economics⁶ that interpret behavior in terms of cognitive costs. A related interdisciplinary psychology literature studying the complexity of (e.g. workplace) tasks (e.g. Donald J. Campbell, 1988; Peng Liu and Zhizhong Li, 2012), makes a useful distinction between “objective complexity” (defined purely by characteristics of a task) and “experienced complexity” (determined also by the abilities, tastes and circumstances of the agent). With this distinction in mind, our experiment is designed to articulate an empirical mapping between “objective complexity” characteristics (e.g. “state complexity”) and “experienced complexity” responses like costs. We will sometimes refer to our cost measures directly as “complexity” (evoking experienced complexity) and sometimes as “the cost of complexity” or “complexity costs” (emphasizing cost as a response to the objective complexity characteristics of the rule).

I. Automata and Complexity

Consider a decision maker (DM) who must make a series of responses ($r_t \in R$) to a series of events ($e_t \in E$) at a sequence of dates $t = 1 \dots T$. We can think of *events* as true historical events (e.g. others’ actions in a game as in Rubinstein (1986)) or as pieces of information considered one at a time in some order (as in e.g. Wilson (2014)). Likewise, we can think of *responses* as either actions made at each date (e.g. actions in a game) or as beliefs developed cumulatively or decisions made provisionally as facts are progressively considered (as in e.g. Salant (2011)).

The DM follows a *rule* that, at each date, maps the history of events so far into a current response. In applications, this could be a rule the DM chose on her own (e.g. which strategy to use in a game or what procedure to use to guide choice) or a rule provided by her environment (e.g. a series of steps and contingencies required to properly file taxes).

The “automata literature” builds on the insight that a rule can be described algorithmically as a *finite state machine* (or *finite automaton*), a basic formalization of algorithms used in computer science (e.g. Hopcroft and Ullman (1979)).⁷ Formally, a *finite state machine* is a four-tuple (S, s^0, f, τ) where S is a set of

⁶There is a recent theoretical literature in economics studying the cost of using complex reasoning procedures (e.g. Haluk Ergin and Todd Sarver, 2010; Pietro Ortoleva, 2013; Larbi Alaoui and Antonio Penta, 2018) as well as a large and growing literature on the costs of attending to information (e.g. Sims, 2003; Caplin, 2016; Dean and Neligh, 2019). Experimental work has been broadly supportive of this latter line of research, producing evidence in favor of the general modeling strategy (Mark Dean and Nathaniel Neligh, 2019) and methods for measuring these attentional costs directly (Andrew Caplin, Dániel Csaba, John Leahy and Oded Nov, 2019; Ambuj Dewam and Nathaniel Neligh, 2020).

⁷Automata can be used to describe a broad range of rules. For instance Kalai and Stanford (1988) show that finite automata can be used to formally describe any finite strategy in a repeated game.

states, s^0 is the initial state, $f : S \rightarrow R$ is an output function, specifying a response in each state, and $\tau : S \times E \rightarrow S$ is a *transition* function that specifies a next state as a function of the current state and the current event. The number of states, $s_i \in S_m$, in state machine m is $|S_m|$, the number of total transitions from state s_i is $N(s_i)$ and the number of total transitions in machine m , summed over states, is $N(m) \equiv \sum_{s_i \in S_m} N(s_i)$. Figure 1 represents two automata visually in the standard way, as directed graphs with (i) circles representing states, (ii) arcs representing total transitions, (iii) letters inside of the circles representing responses from a set $R = \{C, D\}$ and (iv) letters next to transitions representing events from a set $E = \{C, D\}$. A free standing arrow pointing at the left-hand state indicates the initial state, s^0 .

The automata literature poses several hypotheses about what characteristics (e.g. states, transitions etc.) of automata generate *complexity* where “complexity” is interpreted as the cost of *implementing* a rule – a way of ranking the DM’s preferences over implementing different rules. Thus, rule A is more complex than rule B for a DM if the DM would prefer to implement rule B , *ceteris paribus*. We will follow suit by focusing attention throughout the paper on the cost of complexity or “cost complexity.”^{8,9} However, we will also consider other natural metrics of complexity such as implementation time (“rules are more complex if they take longer to implement”) or mistakes rates (“rules are more complex if they are harder to correctly implement”).

A. Structural Hypotheses: States, Transitions and Absorption

We begin with a set of hypotheses that relate to the *structure* of the finite state machine (FSM) representation of the rule. These hypotheses are listed in Table 1 which includes an illustrative example of a relatively simple (on the left) and relatively complex (on the right) rule under each theory.

The earliest (and by far most popular) hypothesis in the automata literature is *state complexity* (or *s-complexity*) (e.g. Rubinstein, 1986; Abreu and Rubinstein, 1988). States describe the distinct ways that a rule requires the decision maker to (i) act (choose a response) and (ii) process information (react to events) and

⁸There is a growing literature in cognitive science developing the idea that using mental resources in complex tasks generates direct costs (e.g. Wouter Kool and Matthew Botvinick, 2018; Amitai Shenhav, Sebastian Musslick, Falk Lieder, Wouter Kool, Thomas L. Griffiths, Jonathan D. Cohen and Matthew M. Botvinick, 2017). A common interpretation in this literature is that complex problems are costly because they require the decision maker to exert “cognitive control” over mental resources that are otherwise deployed reflexively in “automatic” default mental processes. Shutting down these automatic processes (and keeping them at bay) in order to focus cognitive resources on a complex problem that requires dedicated and sustained monopolization of computational resources is costly. Some evidence suggests that this is a top-down function of the executive network in the brain (related especially to the dorsal anterior cingulate cortex). Under this lens, a natural interpretation of our finding that complexity costs fall with learning (Section III.D) is that familiarity allows subjects to follow rules more automatically, requiring them to exert less cognitive control.

⁹Although cost is a natural operationalization of complexity in settings like ours where it is relatively clear that the main thing varying is mental effort, it may raise some conceptual difficulties in other contexts. For instance costs of behavior in some settings may be driven by fear or aesthetic distaste and it may make less sense to refer to this as “complexity.”

Table 1—: Hypotheses on the *structure* of rules.

| Hypothesis | Simple | Complex |
|---|--------|---------|
| s-complexity # of states to track and remember. (Hypothesis 1) | | |
| t-complexity # of events to monitor, respond to. (Hypothesis 2) | | |
| absorption Do tracking and monitoring requirements end? (Hypothesis 3) | | |

Note: On the left we include an example of a simple rule and on the right a complex rule under the Hypothesis listed on the far left.

Table 2—: Hypotheses on the *representation* of rules.

| Hypothesis | Given Form | Reduced Form |
|--|------------|--------------|
| State Reduction Can the rule be represented as a finite state machine with fewer states? (Hypothesis 4) | | |
| Stack Representation Can state-tracking be replaced with a simpler working memory stack? (Hypothesis 5) | | |

Note: Notes: On the left we include an example of the finite state machine that most closely resembles the statement of the rule. On the right we include the simpler way it can be represented under the Hypothesis on the far left.

s-complexity is the claim that partitioning behavior in this way is costly: the DM prefers rules requiring fewer states to those requiring more.

HYPOTHESIS 1: (*s-complexity*) *A rule A is more complex than a rule B if A has more states than B ($|S_A| > |S_B|$).*

Table 1 provides an example of s-complexity: the right hand rule (3S-2T) is more s-complex than the left hand rule (2S-2T) (it has an additional circle in its graph and thus an additional state).

Banks and Sundaram (1990) hypothesize that another central feature of a rule – transitions – drive complexity. Adding a transition to a state (adding additional arcs in the graph of a rule) requires the DM to more precisely monitor and differentiate between events at that state in order to successfully implement the rule. *Transition complexity* (or *t-complexity*) hypothesizes that this increase in monitoring demands is costly. In order to focus on monitoring and to avoid confounding our transitions measure with s-complexity we operationalize t-complexity using “net transitions” (the number of transitions in excess of the number of states) as described in Banks and Sundaram (1990), and we will refer to these simply as “transitions” throughout the paper.¹⁰

HYPOTHESIS 2: (*t-complexity*) *A rule A is more complex than a rule B if A has more transitions in excess of states than B ($N(A) - |S_A| > N(B) - |S_B|$).*

The t-complexity row of Table 1 provides an example: the right hand rule (4S-4T) is no more s-complex than the left hand rule (4S-3T) but is more t-complex because it has one additional arc in its the graph.

Finally, rules can differ in broader “architectural” characteristics, such as whether they contain *absorbing states* – states with no transitions leading away from them. Absorbing rules are potentially less costly and difficult to implement than transient rules (those with no absorbing states) because in the former the DM eventually reaches a position in which she no longer has to track states (s-complexity effectively goes to zero) or monitor the environment (t-complexity goes to zero). Although it is intuitive that “architectural” features of rules such as transience and absorption might influence complexity, the literature does not explicitly model absorbing rules as less complex.

HYPOTHESIS 3: (*transience and absorption*) *Ceteris paribus, rule A is more complex than rule B if B contains an absorbing state while A does not.*

The absorption row of Table 1 gives an example: the rule on the left (4S-3Ta) has identical states and transitions to the rule on the right (4S-3T), but while

¹⁰Each time a reachable state is added to an automaton, at least one additional transition must be mechanically added. Because the goal of the experiment is to differentiate between the effects of elements like states and transitions it is important to use a measure of transitions that does not include this redundancy with our measurement of state-complexity. To do this we use Measure 3 from Banks and Sundaram (1990).

the rule on the left has an absorbing state (the rightmost state of the graph), the rule on the right does not. Under the hypothesis that rules are less complex if absorbing, the rule on the right is more complex.

Two other related ideas from the literature bear mentioning. First, Lipman and Srivastava (1990) point out that automata differ in their *robustness* – in the degree to which mistakes in perceiving events lead to serious mistakes in following the rule’s prescription. We consider this and its bearing on complexity costs in the context of the data in section III.E. Second, Chatterjee and Sabourian (2000) discuss a refinement of s-complexity that they call *response complexity* or “r-complexity,” which applies to settings that differ somewhat from ours (settings in which extensive form games must be played in each state). A rough translation to our setting is that a rule is more r-complex if it requires more changes in behavior (in “responses”) when states change. This is connected to issues discussed in Section I.B below.

B. Representational Hypotheses: Redundancy and Memory

A rule is a way of *representing* a plan of action and, in general, there will be many ways a decision maker can algorithmically represent any plan of action to herself as a rule. Any account of procedural complexity therefore must include an account of how decision makers represent rules to themselves – in particular an account of how *efficiently* they represent rules to themselves. The automata literature has developed under two assumptions about representation that we will use to articulate two *representational* hypotheses in the next two subsections.

ELIMINATING REDUNDANT STATES

First, the automata literature generally assumes that decision makers use representations of rules that economize on states. A rule’s complexity is therefore governed not by the algorithmic features of the rule itself but by the *lowest-state* equivalent rule that delivers an identical sequence of actions over all possible sequences of events. This is an efficiency notion for complexity, postulating that a DM will not waste cognitive effort in following the representation of a rule given to her (or conceived by her) if there is an available equivalent representation that has fewer states. Behaviorally, this is equivalent to a hypothesis that a DM will find a heuristically simpler representation if it is available, re-representing a rule to herself in a less structurally complex way whenever possible.

HYPOTHESIS 4: *(State Reduction) Suppose rule B can be represented as a finite state machine with (i) no more states and transitions and (ii) either strictly fewer states or strictly fewer transitions than rule A. Then rule A is more complex than rule B.*

The State Reduction row of Table 2 provides an example. The rule on the left shows the rule *reducible* which is (under Hypotheses 1-3) identical to the

rule $4S\text{-}4T$ from Table 1. But $4S\text{-}4T$ cannot be reduced to a simpler rule – there is no automaton with fewer states or transitions that delivers the same sequence of actions in response to all histories. By contrast, the rule *reducible* can be replaced with (reduced to) the rule we call *reduced* (shown to the right of *reducible*) – the rule *reduced* produces exactly the same set of actions after all sequences of events as *reducible*, but has half as many states and transitions. Thus, to the degree that states or transitions drive complexity, $4S\text{-}4T$ is more complex in the sense of State Reduction than machine *reducible* even though the two are superficially almost identical. Moreover, to the degree that decision makers represent rules efficiently, we would expect them to treat a rule described as *reducible* as equivalently complex as a rule described as *reduced*.¹¹

REPLACING STATES WITH WORKING MEMORY

More fundamentally, the automata literature implicitly assumes that decision makers represent rules as finite state machines (FSMs) – an assumption that limits the way decision makers can economize on complexity in some cases. For example, the FSM of the rule *countable* (left side of Table 2) has the same number of states, transitions etc. as the FSM of $4S\text{-}4T$ (Table 1), but states serve a much more primitive role in the former rule than in the latter. While states in $4S\text{-}4T$ provide four unique sets of instructions for what to do and how to interpret events, three of the states in *countable* are identical: these states exist purely to keep track of the number of ‘B’ events that have occurred so far. Although states serve a much simpler (largely mnemonic) function for *countable* than $4S\text{-}4T$, under the descriptive lens of FSMs the two rules are identical.

To represent the intuitive difference between the way the two rules make use of states, we must allow the DM to represent rules with algorithmic structures nuanced enough to separate out the memory function of a state from its more complicated functions. This requires only a minimal iteration beyond an FSM called a “pushdown automaton” (PDA). In addition to states and transitions, a PDA maintains a “stack” – a string of symbols – in working memory and can (i) “read” the top element of the stack to determine when to transition, and (ii) can either add (“push”) a new top element to the stack or can remove (“pop”) the top element during any transition. The right side of Table 2, shows a PDA representation of *countable*. Here the transition notation $h, i \rightarrow j$ says “Make this transition if event h has occurred *and* the top element in the stack is i . Replace the top element in the stack with j .” Under this notation, replacing an element with ε removes the element and the element ε is read by the rule only if the stack is empty. The PDA for *countable* begins with an initial transition that

¹¹In the experiment, following the structure of the algorithm, we describe *reducible* as “Choose x until you see b , after which switch to y . Then, choose y until you see a , after which point switch to x . Then, after you see b , switch to y . Then, start over after you see a ” while we describe *reduced* as “Choose x until you see b , after which switch to y . Then, start over after you see a .” Hypothesis 4 argues that subjects should internally replace the former description with the latter.

pushes a string of symbols, “BB,” onto the empty stack ($\varepsilon \rightarrow BB$). The stack is now of length two, ordered from left (“top of the stack”) to right (“bottom of the stack”). In the initial (left hand) state, when a ‘B’ occurs: (i) if the stack is non-empty, the top ‘B’ in the stack is ‘replaced’ with ε and thus removed ($B \rightarrow \varepsilon$) and (ii) once the stack is empty (i.e. once there are no ‘B’ symbols left, meaning the top of the stack contains ε) the automaton transitions to the right.

Representing *countable* as a PDA-like rather than FSM-like rule economizes on s-complexity: it removes two states without adding any transitions, and it does so by instructing the DM to *count* events rather than track many states.^{12,13} Our final hypothesis is that rules that can be simplified in this way, using working memory to replace states, are less complex.

HYPOTHESIS 5: (Stack Representation) *Suppose rule B can be represented as a pushdown automaton with (i) no more states and transitions and (ii) either strictly fewer states or strictly fewer transitions than rule A. Then rule A is more complex than rule B.*

In order to test the hypothesis, we describe the rule *countable* to subjects in the natural way using a PDA-like description: “Choose x until you’ve seen b three times so far (not necessarily in a row), after which switch to y. Then, start over after you see a.” Our hypothesis is that this is less complex than “4S-4T” which can be described in no simpler way than the FSM-like “Choose x until you see a, after which switch to y. Then, after you see b switch to x. Then, after you see another b switch to y. Then, start over after you see a.”¹⁴ Counting can be represented in this particularly simple way in part because it uses an especially simple kind of memory that we call an “event stack”: the rule adds a string of anticipated future events upon entering the initial state, removes elements iteratively as the corresponding events occur and transitions once the stack is empty.¹⁵ Indeed, counting can be thought of as a simple instance of *sequencing*: remembering and eliminating a series of (possibly irregular) events before triggering a change

¹²A natural interpretation is that by separating states from working memory, PDAs distinguish between the (i) “information processing” and “behavioral” functions of a state and (ii) the plausibly less burdensome “memory” function of a state, and thus allow us to attribute different complexity contributions to these functions. In an FSM representation these functions are all constrained to be treated identically.

¹³By contrast, as we show in Online Appendix C, no such “free lunch” is available for rules like 4S-4T: in order to remove states using a PDA representation, the rule must, in exchange, add a great many additional transitions to handle the information processing required by states in that rule.

¹⁴An alternative way to test the hypothesis would have been to describe *countable* to subjects with a much less natural FSM-like description: “Choose x until you see b. Then choose x until you see b. Then choose x until you see b at which point switch to y. Then, start over after you see a.” We strongly suspect subjects would immediately reduce this to PDA form (and wonder why we chose to describe the rule in such an unnatural way).

¹⁵Pushdown automata are more general than finite state automata: any rule that can be represented by the latter can be represented by the former. Indeed, all of the rules we study in our experiment can be represented as 2-state PDAs, generally by adding a great many transitions. However, for most rules this cannot be done with an event stack, and require a far more elaborate coding system on the stack than an event stack allows. In Online Appendix C, we translate all of the rules studied in the paper into two state PDAs and discuss these representational issues in more detail.

in behavior. The rule *sequenceable* at the bottom of Table 2 gives an example. The FSM representation on the left is almost identical to the rule *countable* and like *countable* uses states mostly for memory and therefore can be reduced to a PDA with an event stack (in this case, initialized with “AB”). However, because *sequenceable* requires a richer (irregular) stack, the reduced version of *sequenceable* requires three additional (net) transitions in exchange for the removal of two states, and thus (unlike *countable*) doesn’t strictly fit Hypothesis 5. Nonetheless, depending on the relative complexity costs of states versus transitions, we might expect a net complexity reduction for *sequenceable* as well (though a smaller one than for *countable*).

C. What Else Influences Procedural Complexity?

Other features of rule-following that have little to do with the structure or representation of a rule seem also likely to influence experienced complexity. Perhaps, most important among these is learning and resulting familiarity with a rule. We hypothesize that experience with a rule may allow a decision maker to build up skill that reduces experienced complexity via a process that cognitive scientists call *procedural learning*. Another potentially important factor that may influence complexity, is the context in which the rule is implemented. For instance, some procedures are followed by taking an explicit series of steps (for instance taking a series of actions in a repeated game), while others are followed to form beliefs or to forecast the behavior of others (predicting the consequences of another’s strategies, calculating the consequences of an asset allocation, analyzing a sequence of pieces of information to form a belief etc.). Following a procedure “in one’s head” seems more difficult and therefore plausibly more costly, even if the algorithmic structure of the procedures required in each case are identical.

In the experiment described below we consider both of these potential “non-algorithmic” drivers of complexity. However, we note that there are many other factors, both algorithmic and non-algorithmic, that plausibly influence complexity and that our design does not explore. For instance, it seems plausible that features of the underlying decision problem like the number of potential events or the number of potential responses to events influence complexity. Because none of the hypotheses above concern these features of the decision environment, our experimental design will hold both constant (at two events and two responses). Likewise, none of the hypotheses in the extant literature explicitly concern issues of the architecture of rules (e.g. complex branching structures), beyond those summarized by states and transitions. We stated a hypothesis concerning absorption (Hypothesis 3) which has to do with whether the architecture of automata are “linear” or “circular.” But clearly many, more complicated branching structures could be explored and their influence on complexity characterized. We believe these issues are important topics for future research.¹⁶

¹⁶Likewise, our design hardly exhausts the set of possible influences of *context* on complexity. For



(a) Implementation Task

(b) Elicitation Task

Figure 2. : Screenshots from the experimental software.

II. Experimental Design

The experimental design consists of five Stages (summarized in Figure 3), run in sequence in each session. In Stages 1, 3 and 5 we ask subjects to complete a series of *Implementation Tasks*:

Implementation Task: In each Implementation Task, we assign subjects a verbal *rule* and show them a sequence of twenty ‘Events’ (randomly selected letters), one at a time.^{17,18} The subject’s task is to “implement the rule” by typing the sequence of letters (“Choices”) prescribed by the rule, given the random events that have occurred. Given the sequence of events, there is a unique sequence of twenty Choices that complies with the rule. If the subject perfectly types this sequence, she earns a monetary reward (otherwise she earns nothing). Subjects are not told whether they’ve made a mistake until the end of the Task and are then shown the correct sequence.

Example: *Figure 2 (a) shows an example of an Implementation Task. The subject was shown a rule at the top of the screen (“Choose c until you see g, after*

instance meaning may influence complexity directly. A canonical example is the empirical literature on the Wason selection task (a card task testing deductive reasoning): when the task is framed abstractly subjects do extremely poorly, but they often solve the task correctly when it is framed more meaningfully, for instance as a social exchange or cheater detection problem (e.g. Cosmides, 1987). It is possible that rules implemented in non-abstract, meaningful contexts have lower complexity costs and examining whether this is the case seems a promising direction for future work.

¹⁷For example the rule we call 2S-1T (the Grim rule from Figure 1) has a verbal description “Choose X until you see B, then switch to Y forever,” while the rule 4S-4T (shown in Table 1) has a description “Choose X until you see A, then switch to Y. Afterwards, if you see B switch to X and then to Y if you see B again. Start over if you see A.”

¹⁸These Events are randomly selected by the computer from a set of two possible Events following each Choice, and subjects are informed that Events are randomly selected. At the beginning of each Implementation task, we randomly select the set of two letters used to represent Events and the two possible Choices (responses). That is, we replace the letters in the event set {A, B} and response set {X, Y} each with a randomly selected letter in forming the rule.

which...") and typed a first letter ('c') which was displayed in the first position on the "Choice" line. The computer responded by randomly selecting an Event 'g' (drawn equiprobably from a set of 2 possible letters) and displayed it in red in the first position on the "Event" line. The subject, following the rule, next typed 'l' and a new random Event 'i' immediately appeared in the next position on the Event line, etc. In the example, the subject has made 5 (out of 20 total) Choices in the Task so far.

In Stage 1, we assign subjects 14 Implementation Tasks in sequence, each with a different assigned rule (and each consisting of 20 Choices). The rules are drawn from a fixed set of 14 (described in Section II.A) and are assigned to subjects across Tasks in a random order. At the end of the experiment, the computer draws one Stage 1 Implementation Task at random and the subject earns \$4 if she implemented that Task correctly. The purpose of this Stage is to study how markers of complexity like implementation time (how long it takes the subject to implement the rule) and mistake-making (how likely the subject is to make a mistake) vary with characteristics of the rule.

In Stages 2 and 4 of the experiment we ask subjects to complete several *Elicitation Tasks*:

Elicitation Task: In each Elicitation Task we show subjects a maximally simple "red rule" (directing subjects simply to type the same letter repeatedly) and a more elaborate "blue rule." The subject knows that one of these two rules will be assigned to her in a future Implementation Task. In that future Task, the red rule will pay \$2 while the blue rule (if correctly implemented) will pay some unknown random amount x drawn uniformly from [\$2, \$6]. The subject's task is to use a slider to give the minimum amount m for which she would be willing to be assigned the blue rule rather than the simpler red rule in that future Implementation Task. If $x \geq m$ the subject will be assigned the complex blue rule, otherwise she will be assigned the simple red rule.¹⁹

Example: *Figure 2 (b) shows an example of an Elicitation task. The subject was shown the red rule ("Chose r...") on the left side of the screen and a blue rule ("Choose y until you see w...") on the right. The subject submits m (the "minimum blue payment required") by moving the slider and sees the results ("must pay you at least...") above the blue rule. In this example, the subject has set the slider to \$2.90, meaning she would prefer to be assigned the red rule unless the blue rule pays at least \$2.90.*

¹⁹This task is formally equivalent to a Becker-Degroot-Marschak (BDM) elicitation. Cason and Plott (2014) document that in BDMs framed as auctions, some proportion of subjects confuse the second price auction incentives of the mechanism with first price auction incentives (e.g. they believe they will earn their bid). We designed our mechanism with these results in mind, by (i) explicitly asking subjects to submit an indifference points (rather than presenting an auction), (ii) explicitly highlighting the incentive compatibility of submitting one's true indifference point and (iii) explicitly warning subjects against the specific mistake Cason and Plott (2014) document in their data (i.e. emphasizing to subjects that they are paid the random value x if assigned the complex rule, *not* their indifference point m).

Table 3—: Rules implemented in the design.

| | 2 State | 3 State | 4 State |
|--------------|------------------|--------------------------|---|
| 1 Transition | 2S-1Ta | 3S-1Ta | 4S-1Ta |
| 2 Transition | 2S-2T reduced | 3S-2T | 4S-2T |
| 3 Transition | | 3S-3T 4S-3T 4S-3Ta | |
| 4 Transition | | | 4S-4T reducible countable sequenceable |

The choice, m , gives us a measure of the subject’s subjective cost of implementing the blue rule – the amount the subject is willing to “leave on the table” to avoid implementing the rule in the future. Since the simple red rule pays \$2, our cost measure – the money the subject is willing to sacrifice to avoid the more complex blue rule – is $c = m - 2$ and we measure this cost for each rule assigned in an Elicitation Task. In the example, the subject has submitted \$2.90, meaning her elicited cost, c , of implementing the blue rule is \$0.90.

In Stage 2, we assign subjects 14 Elicitation Tasks, each eliciting willingness-to-pay to avoid a different (blue) rule. These rules are drawn (and assigned in a random order) from the same set of 14 rules we assign for Implementation Tasks in Stage 1, meaning subjects have implemented each of these rules prior to being asked to value them.²⁰ In order to incentivize the elicitations, we inform subjects that at the end of the Stage, one Elicitation will be randomly drawn and used to determine which rule the subject will be assigned to implement repeatedly (in ten separate, sequenced Implementation Tasks) in Stage 3. The purpose of Stage 2 is to measure the complexity cost of each of our rules and study how this cost varies with characteristics of the rules.

A. Rules Studied in the Design

We designed our 14 rules to evaluate the hypotheses in Section I. Table 3 lists all 14 rules, with columns designating the number of states in the finite state machine representation of the rule and rows the number of net transitions (transitions in excess of the number of states). We divide these rules into three types:

- Most cells of Table 3 include a rule with a name of the form “ $xS-yT$ ” where x designates number of states and y number of (net) transitions. These

²⁰After completing the 14 Elicitation Tasks one-at-a-time, subjects are shown their choices for all of the elicitations and given a last chance to fine tune them.

are our **core rules**, designed to test the effects of states and transitions on complexity (Hypotheses 1 and 2).

- Some cells include a rule with a name of the form “ $xS-yTa$,” where ‘a’ designates rules containing an absorbing state. These **absorbing rules** are used to test the effects of transience and absorption (Hypothesis 3).
- Some cells include one or more rule with names that aren’t of the “ $xS-yT$ ” form. These are our **representation rules**, designed to be compared to a core (“ $xS-yT$ ”) rule in the same cell to test Hypotheses 4 and 5. For instance the rules *reducible* and *countable* are designed to be compared to $4S-4T$, with which they share a cell.

Our empirical strategy is to compare behavior across rules in order to measure the average effects of rule characteristics (e.g. states, transitions) on complexity measures gathered in Stage 1 (implementation time, mistakes) and especially Stage 2 (complexity costs). This, in turn, allows us to evaluate our hypotheses. For instance $2S-2T$, $3S-2T$ and $4S-2T$ vary states while holding transitions constant at two. Comparing elicited costs (or implementation time or mistakes) across these rules lets us test s-complexity (Hypothesis 1). Likewise, $4S-2T$, $4S-3T$ and $4S-4T$ vary transitions while holding states constant at four, letting us evaluate t-complexity (Hypothesis 2) by comparing their costs. $4S-3T$ and $4S-3Ta$ share identical states and transitions but only the latter contains an absorbing state, enabling a test of Hypothesis 3. Comparing *reducible* and *countable* each to $4S-4T$ give us tests of Hypotheses 4 and 5 respectively. The other rules are used in the analysis to support and clarify these hypothesis tests.

Most (nine) of the rules from Table 3 are shown in automaton form in Tables 1 and 2 and Figure 1. The remaining rules are shown in Online Appendix A, which also provides the verbal descriptions we gave to subjects.

B. Further Stages

In Stage 3, we assign subjects a single rule to implement repeatedly over ten consecutive Implementation Tasks (each with 20 Events and Choices as in Stage 1). The primary purpose of Stage 3 is to incentivize the elicitations in Stage 2 by assigning subjects rules based on the preferences they submitted. However, a second purpose is to study how learning and familiarity impact complexity. In order to accomplish both goals, we randomly assign half of our subjects as “Endogenous Assignment” types and half as “Exogenous Assignment” types (subjects were never told which type they were). For Endogenous Assignment types, the computer randomly selects one Stage 2 elicitation and a random payment x , and uses these to determine which rule to assign the subject in Stage 3. Exogenous Assignment subjects are instead assigned a rule without reference to their willingness-to-pay, and in the main portion of our dataset, are simply assigned either the rule “ $4S-4T$ ” or “*reducible*.” While Endogenous Assignment subjects

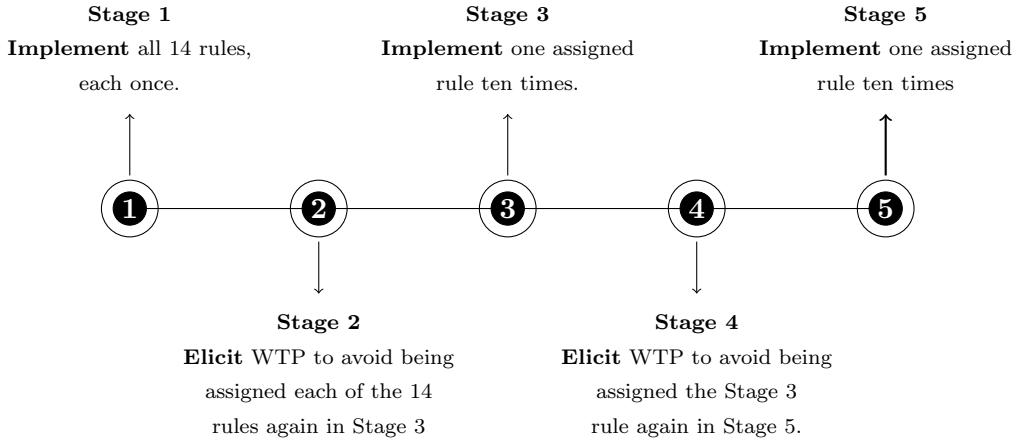


Figure 3. : Summary of the Stages in the experimental design

self-select into their Stage 3 rule via their submitted preferences, Exogenous Assignment subjects do not and their learning data is therefore easier to interpret. Regardless, subjects are paid based on their performance in one of these Implementation Tasks, randomly selected at the end of the experiment.

In Stage 4, we re-elicit subjects' willingness to pay to avoid being assigned (in Stage 5) the same rule they Implemented repeatedly in Stage 3. Our goal with this Elicitation Task is to study whether familiarity with a rule reduces subjects' costs of implementing the rule. By comparing willingness-to-pay before and after learning in Stage 3, we are able to measure the effects of learning on complexity costs for the Stage 3-assigned rule. In order to understand learning effects better, we assign two different Elicitation Tasks in this Stage: one is for the Original version of the rule (the exact same rule implemented in Stage 3) and the other is for a superficially Perturbed version of the same rule. Specifically the Perturbed version is representable by the same underlying automaton as the Original version, but we randomly change all of the letters designating Events and Choices.

In Stage 5, we “pay off” the Stage 4 elicitation by having subjects once again play a rule ten times (ten Implementation Tasks). As in Stage 3, Endogenous Types are assigned based on one of the two Stage 4 Elicitations (randomly selected) and a randomly drawn x , while Exogenous Types are randomly assigned either the Original or Perturbed version of the Stage 3 rule. Subjects are paid based on their performance in one of these Implementation Tasks, randomly selected at the end of the experiment.

C. Details

We ran the experiment using 275 subjects (mostly undergraduates) in March, April, July and August of 2019 at UC Santa Barbara. We recruited subjects using

ORSEE (Greiner (2015)), allowed subjects to participate in only one session each and used custom Javascript software programmed by the author and deployed using Qualtrics. Sessions lasted roughly 90-120 minutes and subjects, on average, earned about \$20.

We ran sessions using three contextual variations on the Implementation task (varied between subjects) which we pool for most of the analysis.²¹ In the **Base** version (110 subjects) subjects follow the Implementation task described above. In the **Recall** version (111 subjects), the task is identical but subjects can only see the *latest* Event and Choice on their screen at any one time. Finally, in the **Reasoning** version (54 subjects) subjects see all 20 events at the very beginning and are asked to forecast what a perfect rule-follower would choose in the very last (of 20) choices, forcing them to implement the rule entirely “in their heads.” By comparing the Base version to the Reasoning version we can study the difference between (i) actually enacting a rule (e.g. implementing a strategy) and (ii) using a rule mentally to make an inference or forecast an event (e.g. reasoning through a decision tree, forecasting another’s choices in a game). On the other hand, while Base mirrors a context in which decision makers have easy access to the history of events (for instance in which ‘events’ are options in a list or pieces of information being procedurally considered in some order), Recall mirrors a condition in which the decision maker must work to remember events (plausibly the case in implementing strategies in some repeated games contexts).

Each session began with detailed instructions on the Implementation task which we read aloud to subjects, and which focused especially on explaining the meaning of terms used in the rules. We encouraged subjects to ask questions during these and during five unpaid practice Implementation Tasks.²² We paused the experiment to verbally deliver instructions prior to Stage 2, and subjects afterwards moved along at their own pace, reading computerized instructions before each Stage. At the end of the experiment subjects participated in a battery of incentivized cognitive tests.²³ Additional details and instructions to subjects are reproduced in Online Appendix D.

We collected data in two “runs” of the experiment: Run 1 (57 subjects under Base, 59 under Recall) and Run 2 (53 in Base, 54 in Reasoning and 52 in Recall). The main difference between the Runs is that we allowed subjects to leave the lab as soon as they finished the experiment in Run 1 but forced subjects to wait until

²¹We pool these variations because, as we discuss in Section III.D, behavior is very similar (except for level shifts easily measured in regressions), across variations.

²²In Reasoning sessions, we gave subjects initial instructions and practice periods from the Base condition in order to help them understand how to read and interpret the rules. We afterwards told them they would be required to forecast final choices rather than fully implement the rules in the actual experiment.

²³We paid subjects \$0.25 for each question they got right. Cognitive tests include a truncated (3-question) version of Raven’s Progressive Matrices (used to test abstract reasoning, fluid intelligence and generalized intelligence), three Belief Bias questions (used to test for logical over intuitive reasoning), three Cognitive Reflection questions (used to measure ability to overcome instinctive responses), the Wason selection task (used to measure deductive reasoning abilities) and five rounds of the Operation Span task (used to measure working memory capacity).

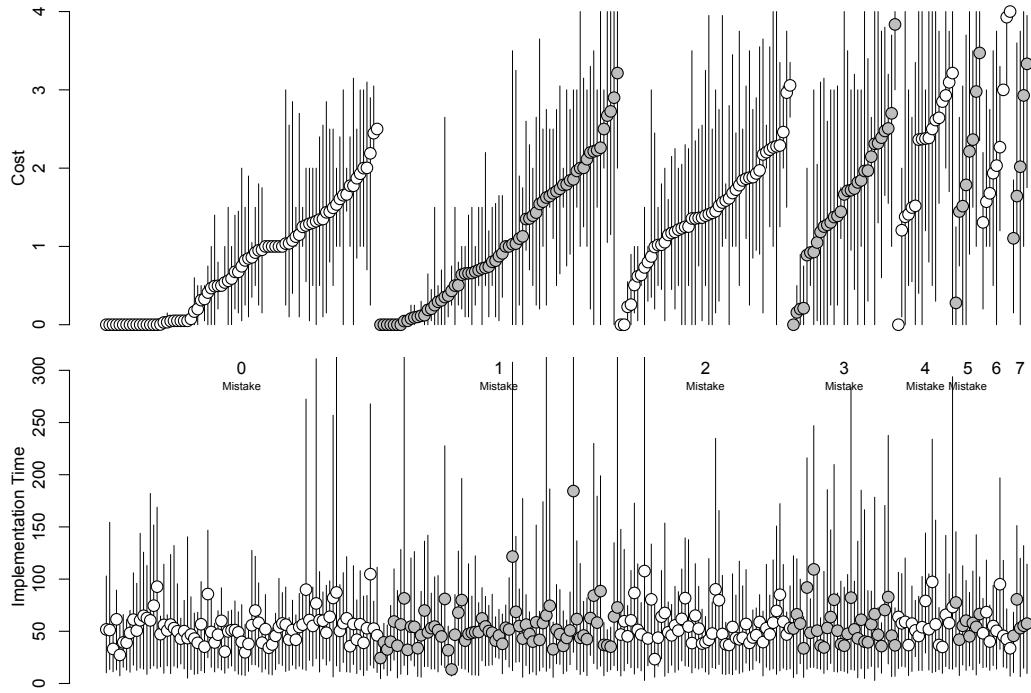


Figure 4. : Average cost and implementation time for all subjects

Source: Subjects are binned according to the number of mistakes they made (out of 14) in Stage 1, and organized from lowest to highest cost within bin. In the top panel the subject's mean cost is shown as a dot and her range of costs across rules as a vertical line. Mean implementation time and range of implementation times is shown in the bottom panel at corresponding horizontal locations for each subject.

everyone was finished in Run 2 (to eliminate the possibility that subjects submit preferences for simpler rules because these rules take less time to implement, allowing them to leave the lab sooner). We also made slight clarifying changes to the wording of a few rules in Run 2 and changed the set of rules we assigned to Exogenous Assignment subjects in Stage 3: in Run 1 we randomly assigned from the full set of rules while in Run 2 we randomly assigned only from the two most measured-complex rules, *4S-4T* and *reducible*.²⁴ The data across the two runs are not statistically different and we pool them in the main analysis below (our results do not change if we instead drop Run 1 from the dataset).

²⁴We made this change to Stage 3 rule assignment in Run 2 in order to better study learning. Run 1 data produced too few observations on any given rule to be statistically useful and rules featured obvious variation in baseline complexity. By assigning subjects only two very complex rules in Run 2 we gather significant data on learning under each rule and are able to focus on a setting in which learning is especially valuable.

III. Results

Figure 4 provides a broad overview of the data at the subject level. The upper panel plots the average elicited complexity cost (over all 14 rules) for each subject in the dataset as a circle. Subjects are organized horizontally by the number of mistakes they made (out of 14) and are ordered from lowest to highest average cost within each of these mistake-count bins. Vertical lines plot the range of costs (from lowest to highest) for each subject. The bottom panel plots the corresponding average implementation time and range of implementation times across rules for each of these subjects (arranged horizontally according to the mistakes/cost ordering from the upper panel). The figure highlights several facts about the data.

First, most subjects make relatively few mistakes: the modal subject makes 0 and the median subject makes 1 mistake over the 14 rules. Second, most subjects – including subjects that rarely make mistakes – are willing to pay a considerable fraction of their expected earnings to avoid implementing the average rule. Just under 10% of subjects are willing to pay nothing to avoid complexity.²⁵ Third, there is a great deal of variation in costs across rules, within subject, indicating that subjects make strong distinctions across rules when evaluating cost: the median least costly rule costs \$0.20 while the median most costly rule costs an order of magnitude more at \$2.35. Fourth, average complexity costs are quite heterogeneous across subjects (the standard deviation of subject-wise means is \$0.91), as are average implementation times (standard deviation 18 seconds). However, at the subject level, there is little relationship between the two: the cross-subject correlation between mean cost and mean implementation time is 0.084 ($p = 0.17$). Finally, regressing costs, implementation time and mistakes rates on a composite statistic of our battery of cognitive tests reveals that all three of these measures are significantly influenced by measures of cognitive ability (see Table 5, described below). That is, for all three measures of experienced complexity, subjects that perform worse in cognitive tasks find rules, on average, more complex.

RESULT 0: *Most subjects make few mistakes but nonetheless find the average rule costly to implement and make strong distinctions in cost across rules. Variation in complexity measures including cost, implementation time and mistakes rates are significantly linked to independent measures of cognitive ability.*

A. Findings on Automata Structure: States, Transitions and Absorption

We now turn to evaluating our hypotheses concerning the determinants of complexity across rules (Results 1-5 below numerically correspond to the numbering of hypotheses in Section I). Figure 5 summarizes the data at the rule level by

²⁵Given our elicitation methods, we cannot rule out the possibility that these subjects experience negative costs and actually prefer complexity.

Table 4—: Rule-level summary statistics.

| Rule | States | Trans. | Absorbing? | Mean Cost | Median Cost | Low Mistake Cost | Mean Time | Mistakes Rate |
|--------------|--------|--------|------------|-----------|-------------|------------------|-----------|---------------|
| 2S-1Ta | 2 | 1 | Y | 0.69 | 0.30 | 0.48 | 15 | 0.01 |
| 3S-1Ta | 3 | 1 | Y | 0.78 | 0.50 | 0.59 | 20 | 0.01 |
| 4S-1Ta | 4 | 1 | Y | 0.89 | 0.65 | 0.64 | 31 | 0.04 |
| reduced | 2 | 2 | N | 1.00 | 1 | 0.71 | 57 | 0.08 |
| 2S-2T | 2 | 2 | N | 1.00 | 0.90 | 0.70 | 60 | 0.17 |
| countable | 4 | 4 | N | 1.23 | 1 | 0.87 | 49 | 0.07 |
| 3S-2T | 3 | 2 | N | 1.29 | 1.05 | 0.93 | 67 | 0.15 |
| 3S-3T | 3 | 3 | N | 1.33 | 1 | 0.91 | 63 | 0.24 |
| 4S-3Ta | 4 | 3 | Y | 1.38 | 1.10 | 0.99 | 35 | 0.04 |
| 4S-2T | 4 | 2 | N | 1.47 | 1.35 | 1.02 | 74 | 0.21 |
| 4S-3T | 4 | 3 | N | 1.59 | 1.50 | 1.15 | 76 | 0.19 |
| sequenceable | 4 | 4 | N | 1.60 | 1.55 | 1.12 | 69 | 0.29 |
| 4S-4T | 4 | 4 | N | 1.73 | 1.80 | 1.28 | 74 | 0.16 |
| reducible | 4 | 4 | N | 1.79 | 1.85 | 1.33 | 73 | 0.09 |

Note: Rules are ordered by the mean cost of the rule. The table lists the number of states in the rule, the number of transitions (net of states) and whether the rule contains an absorbing state. The table also provides mean and median cost, mean time spent implementing the rule and the mistakes rate under the rule. “Low Mistake Cost” is the mean cost for the subset of subjects who made no more than one mistake in Stage 1.

showing average cost, implementation time and mistakes rate as a function of total transitions (transitions + states) across rules. Table 4, shows similar statistics numerically, ordering rules instead by mean cost.

The experimental design allows us to test our main hypotheses about the drivers of complexity by comparing average measured costs (and implementation time and mistakes rates) across rules. For instance, the rules *2S-2T* and *3S-2T* (shown in the *s*-complexity row of Table 1) have the same number of transitions but the latter has one additional state and is therefore more *s-complex*. As Table 4 shows, the average cost for *2S-2T* is 1.00 and for *3S-2T* is 1.29 indicating a $1.29-1.00=0.29$ increase in costs (a 29% increase) from adding a single state to a rule. Averaging across all such controlled comparisons gives us an average effect of adding a state of 0.236 (a 20% increase), supporting the hypothesis that states generate complexity costs (our Hypothesis 1).

Formally, we test our hypotheses using a set of regressions, summarized in Table 5. Each specification corresponds to a different complexity metric (cost, implementation time and mistakes rate) and includes right-hand variables for the number of states, the number of transitions and an indicator variable for whether the rule is absorbing (interacted with both state and transition counts). The regressions also include indicators for the rules *reducible*, *countable* and *sequenceable* which we will examine in the next subsection, indicators for the Run of the experiment and for the context (Reasoning, Recall) in which the rule was implemented, and some basic demographic variables. Standard errors are clustered at

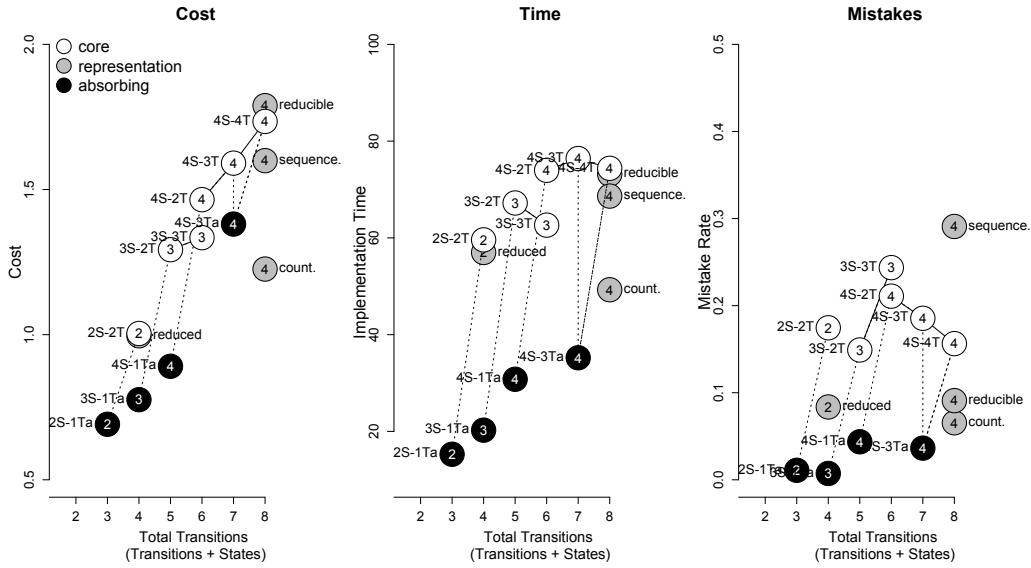


Figure 5. : Average complexity measures for all rules.

Source: Each circle in each panel represents a different rule. Horizontal location of each rule is “Total Transitions” (states plus transitions) and the number of states are shown inside each rule’s circle. Lines connect all of the non-representation rules that contain the same number of states; dotted lines mark transitions that cause the rule to become transient (non-absorbing). Panels show mean cost, implementation time and mistakes rates respectively.

the subject level.²⁶ Estimates from these regressions are visualized in Figure 8.²⁷

We begin by examining the non-absorbing “core rules” ($2S-2T$, $3S-2T$, $3S-3T$, $4S-2T$, $4S-3T$ and $4S-4T$), which were designed to be compared directly with one another to study state and transition effects (s-complexity and t-complexity). The ordering in Table 4 (which, recall, is by mean cost), clearly shows that these rules are ordered first by state (every 4 state rule is more costly than every 3 state rule in this set, etc.) and secondarily by transitions. Moreover, the six most complex rules in the dataset are those with the maximal number of states used in the design. Table 5 confirms that, on average, adding a state has a large, significant impact on cost (equal to almost 20% of the average cost and a 26% increase over the intercept). States also significantly impact implementation time

²⁶A caution in interpreting the standard errors from these estimates is that some effects (e.g. states, transitions, absorption) are estimated using multiple rules and therefore multiple decisions by each subject, while others (e.g. countability, reducibility) are estimated using only a single per-subject choice.

²⁷In Online Appendix B we include several robustness checks on our main cost specification including a Tobit and a specification that includes the word count of the rule as a dependent variable. In neither case do the results of our main Hypotheses tests change. For the Tobit, estimates are similar though we lose precision and significance on the *sequenceable* indicator. For the word count specification, the word count variable is not significant but other variables are attenuated due to its inclusion; the main qualitative change is that the *sequenceable* indicator, again, ceases to be significantly different from zero at conventional levels.

Table 5—: Estimates of drivers of complexity metrics.

| | Cost (1) | Time (2) | Mistakes (3) |
|---|-------------------|--------------------|-------------------|
| <i>States</i> (H1 , # of states in rule -2) | 0.239 (0.021) | 8.788 (1.083) | 0.027 (0.012) |
| <i>Transitions</i> (H2 , # of transitions in rule -1) | 0.116 (0.019) | -0.752 (1.357) | -0.003 (0.013) |
| <i>Absorbing?</i> (H3 , indicator absorbing state in rule) | -0.203 (0.039) | -44.098 (1.959) | -0.142 (0.020) |
| <i>Reducible?</i> (H4 , indicator for rule ‘reducible’) | 0.072 (0.034) | -0.621 (2.527) | (0.025) |
| <i>Countable?</i> (H5 , indicator for rule ‘countable’) | -0.490 (0.048) | -24.491 (2.254) | (0.024) |
| <i>Sequenceable?</i> (H5 , indicator for rule ‘sequenceable’) | -0.117 (0.041) | -5.128 (2.371) | (0.030) |
| <i>Recall?</i> (indicator for Recall sessions) | 0.020 (0.114) | 0.359 (1.548) | 0.005 (0.017) |
| <i>Reasoning?</i> (indicator for Reasoning sessions) | 0.301 (0.150) | 17.473 (4.147) | 0.015 (0.020) |
| <i>Run 1?</i> (indicator for Run 1 of the experiment) | 0.047 (0.112) | -3.062 (1.582) | 0.006 (0.018) |
| <i>Female?</i> (indicator for gender, 1 = Female) | -0.021 (0.105) | -1.047 (2.035) | -0.016 (0.016) |
| <i>STEM?</i> (indicator for STEM major) | -0.157 (0.106) | 0.334 (2.340) | -0.006 (0.015) |
| <i>Cog Score</i> (cognitive battery score) | -1.801 (0.337) | -11.445 (6.897) | -0.287 (0.048) |
| <i>Absorbing X State</i> | -0.140 (0.023) | -1.037 (1.276) | -0.010 (0.014) |
| <i>Absorbing X Transition</i> | 0.131 (0.027) | 3.435 (1.456) | 0.003 (0.014) |
| Constant | 0.903 (0.123) | 56.560 (2.973) | 0.151 (0.027) |
| Observations | | 3,850 | 3,850 |
| R ² | | 0.182 | 0.337 |
| | | | 0.081 |

Note: The dependent variables include cost (specification 1), implementation time (specification 2) and an indicator taking a value of 1 if the subject made a mistake (specification 3). Each specification clusters standard errors at the subject level. Right hand variables include number of states and transitions and a dummy for whether the rule includes an absorbing state. We subtract 2 from states and 1 from transitions (i.e. the minimal number of states and transitions in the design) in order to make the intercept easier to interpret. Reducible?, Countable? and Sequenceable? are dummy variables corresponding to the representation rules of the same name. Recall and Reasoning are dummies for the versions of the Implementation Task of the same name. Run 1 is a dummy variable for the first run of the experiment and Female and Stem are gender and college major dummies. Cognitive Battery is an index of scores from a battery of cognitive tests, ranging in value from a minimum of 0 to a maximum of 1 and demeaned.

and mistakes rates.

RESULT 1: *Adding a state to a rule has a large and significant positive effect on the cost of implementing it.*

By comparison, Table 5 shows that adding a transition, though statistically significant, adds half as much cost as adding a state.²⁸ Transitions have no significant effect on implementation time or mistakes.

RESULT 2: *Adding a transition to a rule has a significant positive effect on its cost, but the effect is half as large as that of adding a state.*

The Absorbing indicator in Table 5 is large and statistically significant: changing a rule from absorbing to non-absorbing adds as much cost as adding a state. Indeed, as Table 4 and Figure 5 show, the three least costly rules are all absorbing rules. Absorption also blunts the impact of states on complexity costs: adding states to absorbing rules has a much smaller impact on complexity costs than does adding states to non-absorbing rules (*Absorbing* \times *State* is significantly negative). Finally, because all transitions added to absorbing rules simultaneously remove the absorbing state in our design, the large significant effect of the interaction *Absorbing* \times *Transition* tells us that adding a transition that simultaneously causes an absorbing state to become non-absorbing, has double the effect of adding a conventional transition. Absorption has a similarly sizable impact on our other complexity measures.

RESULT 3: *Absorbing rules are significantly less costly than non-absorbing rules, equivalent to removing a state. Adding a state to an absorbing rule generates substantially less cost than adding one to a non-absorbing rule.*

To summarize, structural automaton characteristics collectively have strong effects on complexity costs: at the median, the most complex rule (according to Hypotheses 1-3) is six times more costly than the least complex rule. States and absorption (but not transitions) have qualitatively similar effects on implementation time and mistakes rates, our other two complexity measures.

B. Findings on Representation

We next examine the hypothesis (Hypothesis 4) that decision makers efficiently remove unnecessary states from rules by representing them in their least state-complex form. Recall that the rule *reducible* has identical states, transitions and transience to rule *4S-4T* but can be represented as (reduced to) the much simpler (2 state, 2 transition) rule *reduced*. We find no evidence that subjects do this: Figure 5 shows that the elicited cost of implementing rule *reducible* is in fact slightly *higher* than its state/transition equivalent *4S-4T* (and the *reducible*

²⁸There is some evidence of diminishing marginal costs in states but not in transitions. The difference between *2S-2T* and *3S-2T* is larger than the difference between *3S-2T* and *4S-2T* ($p < 0.01$, paired Wilcoxon test), but the difference between *4S-2T* and *4S-3T* is nearly identical to the difference between *4S-3T* and *4S-4T* ($p = 0.947$).

dummy is statistically significant and positive in our cost specification).²⁹ Furthermore, as is clear from Figure 5 and Table 4, *reducible* is far more costly than the rule it can be reduced to, *reduced*.

RESULT 4: *Subjects do not economize on state complexity by reducing rules to simpler isomorphic finite state machine representations.*

Finally, we test the hypothesis that decision makers are constrained to remember events by tracking states, representing rules exclusively as finite state machines (FSMs). Our alternative hypothesis (Hypothesis 5) is that when it is possible to replace states with simple sequences of events in working memory, subjects do so and that this reduces complexity. Our key rule for examining this hypothesis is *countable* which is identical to the rule *4S-4T* when represented as an FSM, but can be represented as a 2-state, 4-transition rule if represented as a pushdown automaton (PDA) that can track elapsed events in working memory.

Figure 5 strongly supports Hypothesis 5: subjects treat the rule *countable* as substantially less complex than *4S-4T*. The *countable* variable in the cost specification of Table 5 is highly significant and quantitatively large. Indeed, the magnitude of this reduction is highly consistent with what we would expect if subjects represented the rule as a PDA rather than an FSM: removing 2 states and 0 transition from the rule (as reduction from a 4-state, 4-transition FSM to a 2-state, 4-transition PDA would dictate) should lead to a 0.472 reduction in cost (based on estimates from Table 5), which is nearly identical to our -0.49 estimate. Thus the cost reduction we observe for *countable* is exactly what we would expect if subjects used a PDA-like representation rather than an FSM-like representation. (Implementation time and mistakes rates are also significantly smaller for *countable* than for *4S-4T*.)

RESULT 5: *Rules that can be represented with fewer states (and no more transitions) by using working memory to track events are significantly less complex.*

A similar but (as expected) weaker pattern occurs with the rule *sequenceable* which can also be state-reduced using a PDA-like representation, though only by adding transitions in exchange for the removed states. Again, Figure 5 shows (and the *sequenceable* dummy in Table 5 confirms) that complexity costs and implementation time are smaller for this rule than for the reference rule *4S-4T*. The cost drop is much smaller than the drop for *countable* (0.11 vs. 0.49) but as we argue in Section I.B, we expect the effect to be much smaller here because the PDA

²⁹This is true, despite the fact that subjects make half as many mistakes in *reducible* as in *4S-4T*. This low mistake rate is likely a consequence of the fact that the pattern of actions required by *reducible* is relatively robust to errors in correctly perceiving and responding to events (Lipman and Srivastava, 1990), a characteristic that is orthogonal to the fact that the rule is reducible. See Section III.E below for a discussion of this.

representation requires several more transitions than the PDA representation of *countable*.³⁰

In Online Appendix C, we consider the merits of describing rules as pushdown automata rather than finite state machines in more depth. Our findings suggest that further study of what classes of algorithms (e.g. finite state machines, pushdown automata, Turing machines etc.) best describe and predict complexity is an important direction for future research.

C. Heterogeneity and Prevalence

We next consider heterogeneity of these effects across subjects. First, recall (Figure 4), that a cluster of subjects (about 17%) submit identical WTP numbers across *all* of the rules. Most of these are subjects who never adjusted their sliders at all, causing them to submit a default cost of 0 perhaps because they suffer no complexity costs (or are complexity seeking). For the remainder of subjects, we individually fit our main estimating equation from Table 5 for the determinants of cost and implementation time, and plot our key coefficients as scatter plots (one for each dependent variable) in Figure 6. We highlight several points.

First, as kernel density plots of the marginal distributions, plotted on the edges of these graphs, show, some of the cost effects we've documented occur for almost all of these subjects (almost 90% for s-complexity and countability), some occur for most of them (2/3 for transitions and absorption) and some for only half of them (reducibility and sequenceability). Second, as the scatter plots show, the effects of rule characteristics on complexity costs and implementation time are generally uncorrelated within-subject or only very weakly correlated: in none of these cases is there a correlation coefficient greater than 0.16 and all but absorption are statistically insignificant. Thus knowing the size of the effect of e.g. states on implementation time for a subject provides little information on the size of its effect on the subject's complexity cost.

Finally, there is significant heterogeneity in the magnitude of each effect across subjects, meaning there is considerable unexplained subject-to-subject variation in *quantitative* effects on costs and implementation times in the data. We can quantify the heterogeneity of these effects by studying partial correlations between our main independent variables and our complexity metrics (cost and time).³¹ The results of this analysis show, for example, that although states

³⁰The estimated cost reduction for *sequenceable* is, again, exactly what we would expect, given our estimates, from removing 2 states and adding 3 transitions from the rule (as would occur by transitioning to a PDA representation). At estimated costs of states for transitions, we should expect *sequenceable* to be about one transition less costly than a rule represented as a 4 state, 4 transition FSM and that is exactly the size of the *sequenceable* estimate.

³¹The partial correlation of y and x is calculated by regressing each on a vector of other relevant variables \mathbf{z} and then calculating the correlation coefficient between the residuals. It describes how much of the variation in y is explained by x after controlling for the explanatory contribution of other relevant variables. We calculate partial correlations on subject-wise demeaned costs and include variables for all of our hypotheses by including state and transition counts, an absorption indicator and indicators for the *countable*, *sequenceable* and *reducible* rules. Since states and transitions in particular have very different

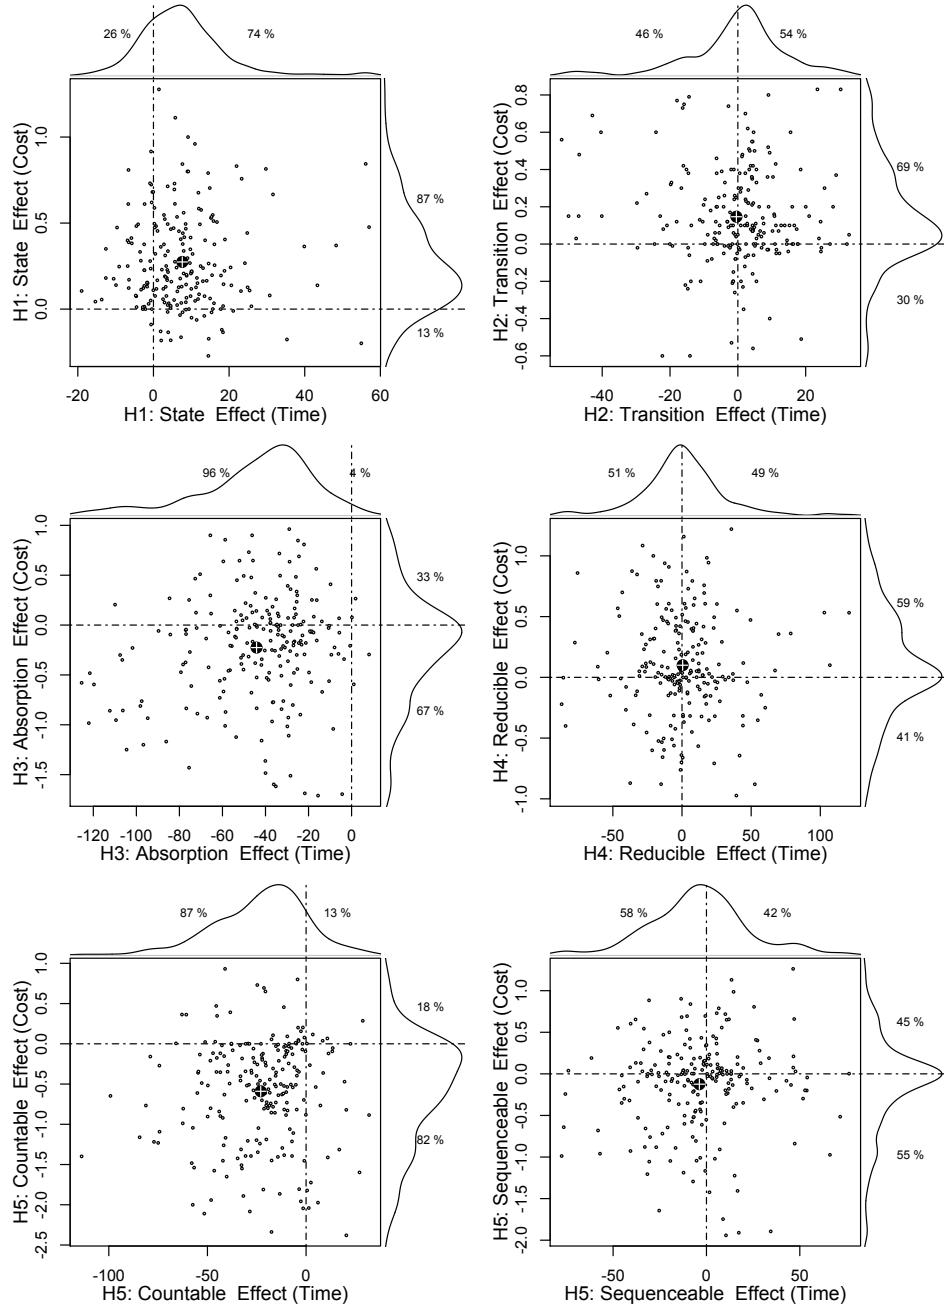


Figure 6. : Individual regression effects.

Note: Each panel shows a scatter plot (with implementation time on the x-axis and cost on the y-axis) of estimates from individual subject-level regressions using specifications (1) and (2) from Table 5. Labels give the name of the variable reported in the panel (matching variable names in Table 5) and H1,...H5 labels give the corresponding hypothesis numbers. On the edges of the graphs are plotted kernel density estimates of the marginal distributions of estimates for implementation time (top edge) and cost (right edge). Numbers on each side of zero on these density plots report the percentage of subjects with negative vs. positive estimated coefficients for the variable. We cut out 2.5% outliers from either side for scatterplot visibility.

increase complexity costs for the great majority of subjects, the magnitude of the effect is heterogeneous enough that they predict only about 25% of residual variation in costs even after demeaning cost at the subject level. Smaller fractions of the variance are explained by countability (22%), absorption (14%), transitions (11%) and reducibility (3%). Corresponding numbers for implementation time include absorption (47%), countability (21%), states (18%), transitions (1%) and reducibility (< 1%). Thus, although many of these characteristics impact complexity in a consistent direction for most subjects and have quantitatively large effects on average, the magnitude of their effects on individual responses are quite varied.

D. Learning, Familiarity and Context

Stages 3-5 from Run 2 of the experiment allow us to study the influence of *familiarity* on complexity directly by allowing subjects to procedurally learn rules via repetition. In Stage 3 we randomized half of our subjects into one of two maximally complex rules: *4S-4T* and *reducible* (the two most cost-complex rules in our data). These subjects then implemented their assigned rule 10 times (each time in response to 20 events). Panel (a) of Figure 7 shows the evolution of mean implementation time over the 10 tasks of Stage 3. Implementation time falls substantially with experience, eventually falling to 66% of initial levels (Wilcoxon test, $p < 0.001$, $N = 84$).

Familiarity also reduces the complexity costs of implementing a rule. In Stage 4, we elicited subjects' willingness to pay to avoid *implementing the same rule they had previously implemented repeatedly in Stage 3* for a further 10 tasks (in Stage 5). In order to focus on the effect of learning and familiarity on preferences (rather than on the rate of mistakes), we restrict attention to subjects who made almost no mistakes to begin with in Stage 1.³² The datapoint "Learn" in panel (b) of Figure 7 presents the change in cost between initial elicitation for this rule in Stage 2 and the re-elicitation in Stage 4. The effect is statistically significant (paired Wilcoxon test, $p < 0.001$, $N = 42$) and quite large, equivalent to stripping more than two states from the rule.³³

These familiarity and procedural learning effects are strong but they are also quite superficial and specific to details of the exact rule learned. In Stage 4, we also elicited costs for a very slight perturbation of the rule implemented in Stage 3. In the perturbation, we replaced each possible choice and event in the rule

characteristics under transient than absorbing rules, we calculated these and other non-absorbing effects using only data from transient rules.

³²The results are similar but attenuated if we use the entire sample instead: using the entire sample, the effect of learning on costs for the original rule (perturbed rule) is -0.371 (-0.169) instead of -0.479 (-0.25). Wilcoxon tests indicate an effect significantly different from zero in all cases.

³³Subjects (in Stage 3) learn faster under the rule *reducible* than under rule *4S-4T*, dropping 37 seconds rather than 22 seconds ($p = 0.003$, Wilcoxon test), but there is no difference in cost reduction subjects achieve with experience across these two rules ($p = 0.731$). This may suggest that, at least on one complexity dimension, some subjects learn to reduce rules (Hypothesis 4) to some degree with experience.

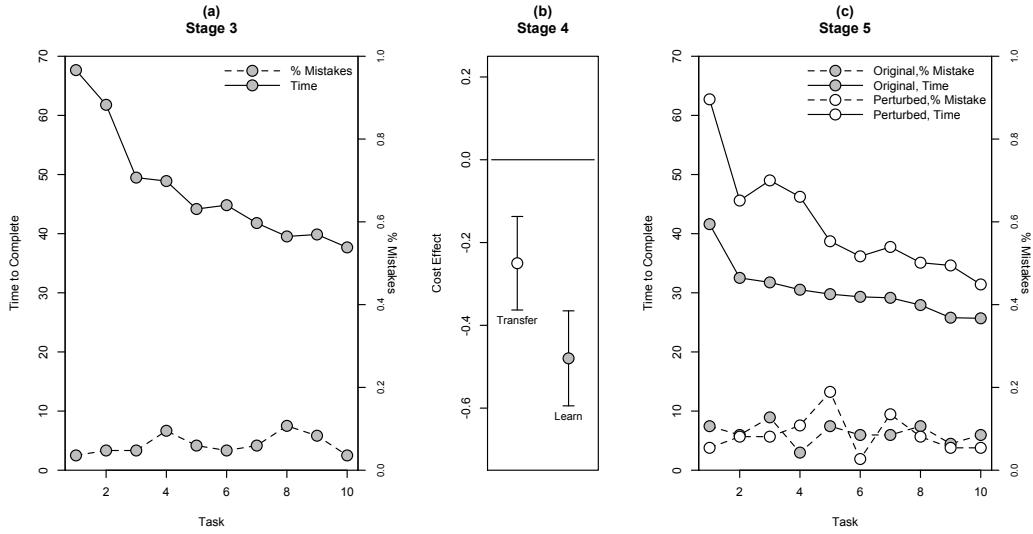


Figure 7. : Learning and familiarity effects.

Note: Panels for Stage 3 and Stage 5 present mistakes rates (dashed lines) and mean implementation times (solid lines) over 10 Implementation Tasks (each task is a separate Implementation Task consisting of 20 events). The panel for Stage 4 shows the mean difference (with standard error bars) between Stage 2 cost elicitation and Stage 4 cost elicitation for the rule implemented in Stage 3 (Learn) and a perturbation of this rule (Transfer).

(for example the letters ‘a’ and ‘x’) with a random set of letters (for example ‘a’ is replaced with ‘b’ and ‘x’ is replaced with ‘y’ throughout). Structurally (and under the lens of our Hypotheses) this “perturbed rule” is identical but it has been superficially – almost cosmetically – altered. The datapoint “Transfer” in panel (b) of Figure 7 presents the change in cost (relative to the original elicitation of the non-perturbed rule in Stage 2) using the exact same subjects shown for the datapoint “Learn.” There is still a significant reduction in cost (paired Wilcoxon test, $p = 0.025$, $N = 42$), but the size of the effect is half that observed for the original rule ($p < 0.001$, $N = 42$).

In Stage 5, we similarly examined whether the effects of learning on *completion time* transfer to superficial perturbations of the learned rule by randomly assigning some subjects to implement the original rule and another set to implement the perturbed rule for 10 tasks. In panel (c) of Figure 7 we plot a time series of implementation times for both the original rule and its perturbation. While the implementation time of the original rule continues its Stage 3 trajectory downward with only a slight restart effect, *nearly all of the implementation time reduction* achieved in Stage 3 disappears for the perturbed rule: subjects must start the learning process almost entirely over again (see Cooper and Kagel (2010) for a review of related findings on “learning transfer” in prior experiments).

RESULT 6: *Familiarity resulting from procedural learning substantially reduces complexity – the effect is equivalent to removing two states from the rule. However, this effect is extremely rule-specific and the impact of familiarity on complexity does not fully transfer to even slight perturbations of the learned rule.*

The *context* in which a rule is implemented also has effects on complexity. In our design we varied the context of our Implementation Tasks (without changing the algorithmic structure of the rules assigned to subjects) in such a way as to vary whether subjects had to perform the rule mentally (the Reasoning variation) or remember past events (Recall) to implement the rule. The estimates in Table 5 include indicator variables Recall and Reasoning to measure main effects of these variations. The results reveal that subjects find implementing a rule mentally to form a belief about the correct final choice prescribed by the rule (the Reasoning variation) considerably more costly than implementing the full rule via a sequence of actual actions (Base). The size of the effect is large, equivalent to adding more than one state to the rule, on average. Reasoning through a rule is also considerably slower but does not produce more mistakes. By contrast, having to remember recent history to perform the rule (as subjects must do in the Recall version) seems to have no significant effect on any measures of complexity.

RESULT 7: *Reasoning through a rule is considerably more costly and time-intensive than enacting a rule. The effect is equivalent to adding more than one state to the rule, on average.*

Interestingly, although the Reasoning variation produces an increase in average cost, it does not much alter the way automata characteristics influence these costs. Indeed, interacting indicator variables for “Recall” or “Reasoning” with the regressors in Table 5 produces no significant interactions and plotting (a’la Figure 5) costs separately for each variation shows essentially the same cost function (with respect to e.g. states or transitions) across variations. (The only exception is absorption: costs across variations are similar for absorbing rules, but non-absorbing rules are are much costlier under Reasoning than under the other variations.) Our results thus suggest that the algorithmic drivers of complexity costs we identify are remarkably robust to variations in context.

E. Complexity Costs, Implementation Time and Mistakes

How does complexity cost relate to other metrics of complexity, particularly implementation time (a common complexity metric in economics experiments, see e.g. Gill and Prowse (2018))? Figure 8 summarizes our findings by plotting the effects (as a proportion of the average) of hypothesized drivers of complexity on costs vs. time (most from estimates in Table 5). Clearly cost and implementation time are shaped by similar forces: the drivers of each are strikingly similar in Figure 8, with proportional effects that are mostly linearly related on the 45-degree line. However, as individual level data visualized in Figure 4 and 6 make clear,

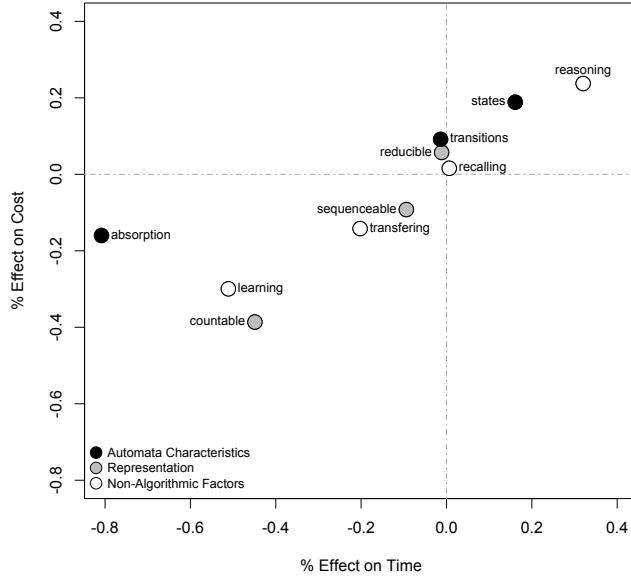


Figure 8. : Estimates of effects on cost complexity and implementation time.

Note: Most estimates are taken from (and match in name) coefficient estimates from Table 5, and are reported as a fraction of mean cost/implementation time. Learning and transferring measures are taken from treatment effects in Section III.D.

the two are actually only weakly related at the individual level: both individual averages and coefficients are insignificantly correlated across these measures. Indeed including time in specification (1) of Table 5 (see Online Appendix B) reveals only very weak explanatory power of a rule's implementation time on cost at the subject level. We conclude that although complexity costs and implementation time are driven in the aggregate by very similar forces (both quantitatively and qualitatively), these are relatively independent effects, underscoring the likely importance of measuring complexity costs directly.³⁴

By contrast, even in the aggregate, the drivers of costs and *mistakes* are only very weakly related, as Figure 5 and Table 4 show. For instance, the rule *reducible* and its equivalent *reduced* have wildly different costs and implementation times attached to them but they are almost identical in terms of mistakes. The reason for this is that, as Lipman and Srivastava (1990) show, automata vary dramati-

³⁴This finding is surprising because it is intuitive to think of complexity costs as arising from per-second costs of time spent enacting a rule. One possibility is that there is a strong latent relationship between time and cost that is attenuated by noisy measurement of each in our experiment. Another possibility is that subjects are averse to complex rules in large part because of fixed cognitive costs (e.g. the costs of absorbing and thinking through the structure of the rule) that have relatively little to do with the length of time they will have to spend implementing the rule.

cally in their *robustness* to mistakes in accurately tracking the sequence of events and *reducible* and *reduced* are identically robust in this sense. By contrast, the rules 2S-2T and *reduced* have identical states and transitions and are treated as equally costly by subjects, but the former is far more forgiving of mistakes in tracking. As a result, subjects make mistakes twice as often in 2S-2T as they do in *reduced* (and in the equivalently robust rule *reducible*). Likewise, subjects treat the rule *sequenceable* (which can be represented with fewer states using a memory stack) as less costly than similar state/transition rules (e.g. 4S-4T) but make mistakes twice as often. Subjects clearly do not take mistakes into account in a sophisticated way when assessing cost (even after having experienced each of these rules!), and so do not attach low costs to low-mistake rules like *reduced* and *reducible* or high costs to high-mistake rules like *sequenceable*.

This lack of association between costs and mistakes holds more generally and this is important for interpreting our cost estimates. *Ex ante*, WTP to avoid complexity might have been driven either by (i) sheer subjective distaste for complexity or (ii) fear of making future mistakes (and thereby sacrificing expected earnings). In Online Appendix B we provide evidence that, in fact, the drivers of complexity costs we have identified are overwhelmingly driven by the former (distaste for complexity). Repeating our cost analysis on subjects who make very few mistakes (and therefore are unlikely to pay to avoid mistakes in the elicitation) or including mistakes explicitly as a dependent variable generates almost identical parameter estimates to those reported in specification (1) of Table 5. This, in turn, suggests that most of the cost-generating effects (e.g. s-complexity, t-complexity etc.) measured in the experiment are driven by subjects' distaste for implementing complex features of rules.³⁵

Our design purposely minimizes the relationships among these complexity metrics, but in many realistic settings they are likely to be more entangled. For instance our design minimizes the opportunity cost of implementation time, but in many natural settings we expect such costs to contribute significantly to the cost of complexity. Likewise, our subjects are not time constrained, but in many realistic settings time constraints will surely produce a higher rate of mistakes for complex rules. Finally, our incentives are strong enough to motivate subjects to (mostly) properly implement rules, but in many settings this won't be true (either because of higher complexity costs or lower incentives), leading to a greater incidence of apparent "mistakes."

³⁵Mistakes rates are low enough in our experiment that it is relatively easy to separate the contribution of distaste for complexity from fear of making mistakes. Extending this research strategy to more difficult tasks may require iterations on these methods that separately elicit (i) beliefs about the likelihood of mistakes and (ii) willingness to pay to avoid the task in the future, making it possible to separate distaste for the task from fear of payoff consequences from making mistakes directly.

IV. Applications

In the introduction we emphasized that procedural complexity has potential applications across economics, and we close the paper by illustrating how evidence like ours might be useful for interpreting and predicting behavior in applications. To do this, we fit our complexity cost estimates to the automaton characteristics (number of states, transitions etc.) of some important rules in economics (rational choice procedures and repeated game strategies), and relate these calibrations to observations from prior empirical work. These are, of course, *ex post* exercises but our purpose in doing this is primarily to hint at the potential uses of our methods for future, prospective research.

A. Rational Choice Procedures

Rational choice is a core assumption in economics, but there is a long tradition beginning with Simon (1955) arguing that rational choice is too procedurally complex to be a realistic description of human behavior. When choosing between options (e.g. in a list), it is much simpler to instead *satisfice* by setting a reservation value (“aspiration level”) and choosing the first option that meets or exceeds it (instead of exhaustively considering all options as is required by rational choice). Salant (2011) uses an automata model to show that even in very basic settings rational choice is far more procedurally complex than satisficing: procedures that select the optimal (most valuable) option from a list of N options are extremely s -complex, requiring $N - 1$ states, while satisficing requires only 1 state regardless of N . Satisficing and related simplified choice procedures are simple, but they lead to suboptimal choices, producing classic behavioral phenomena like recency and primacy effects, status quo bias and choice overload.

We calibrate the cost of rational choice using Salant’s characterization for lists containing $N = 10, 20$ and 40 options and plot the results in the left panel of Figure 9. Salant focuses on the consequences of assuming that states drive complexity: following this and using our estimated cost of states (and including the penalty from implementing a rule mentally reported in Result 7 and the intercept) we get complexity cost estimates of using rational choice procedures that range from just under \$3 for 10 option lists to around \$10 for 40 option lists. Moreover, our results suggest this is probably an underestimate: if we add the cost of the $(N+1)N/2$ transitions required to implement rational choice, the complexity cost rises to as high as nearly \$80 for a 40 option list!

How might we use complexity calibrations like these to predict and interpret behavior? We illustrate using an experimental design reported in Caplin, Dean and Martin (2011) in which subjects evaluate and compare options from a list and make a choice after consideration. Their setting is useful for this exercise because the options their subjects choose between have objective monetary values: subjects are shown the monetary payoff attached to each object in an opaque fashion (as a list of numbers to be added together) meaning some effort and

memory is required to evaluate and compare. Caplin, Dean and Martin (2011) vary the size of their option lists between 10, 20 and 40 in their design and we can use the lists employed in the experiment to calculate, for comparison, the opportunity cost of using a satisficing procedure (the expected money sacrificed by satisficing instead of using rational choice) for a range of possible satisficing reservation values (satisficing has one state and so a 0 cost in our framework). We plot these expected satisficing costs for reservation values 5, 10, 15, 20 and 25, shown from left to right as a dotted line (for each list length) on the plot.

Even for the smallest lists, the complexity costs of rational choice dwarf the expected payoffs sacrificed by satisficing and so we would expect satisficing to be common (and rational choice rare) in all three cases based on procedural complexity alone. Indeed, Caplin, Dean and Martin (2011) find strong evidence of satisficing for all three list sizes, with subjects regularly choosing payoff dominated options and, further, displaying distinctive order effects that indicate satisficing (i.e. when the highest value item is lower in the list, subjects make more mistakes). Moreover, choice process data that tracks subjects' provisional choices show distinctive evidence of satisficing.

Our calibration suggests that rational choice may be procedurally complex enough to induce high rates of satisficing, but we emphasize that careful prospective experiments are required to convincingly identify procedural complexity as an explanation and especially to separate its role from that of other candidate explanations (such as e.g. optimal search behavior).³⁶ For example, experiments that correlate individual procedural complexity costs with satisficing or that systematically vary the opportunity cost of satisficing relative to the complexity costs of rational choice would provide much stronger evidence. Developing synthetic choice environments that vary procedural complexity costs without altering the predictions of alternative models like rational search would be particularly valuable (see Footnote 40, below, for an example). More generally, using these research strategies to study the capacity of procedural complexity to explain and tie together other key regularities from behavioral economics such as random choice (Kalai and Solan, 2003), non-Bayesian inference (Chauvin, 2020), biases in information processing (Wilson, 2014), failures of backwards induction (Neyman, 1985) and perhaps others such as myopia in intertemporal choice, failures to contingently reason through decision trees, status quo bias etc. is an important next step in this research.

This exercise also illustrates some of the subtleties involved in deploying automata models in practice and points towards potentially important extensions of these models. One important finding from Caplin, Dean and Martin (2011)

³⁶Caplin, Dean and Martin (2011) find that comparative statics from search models incompletely organize differences in behavior between their treatments, concluding "our subjects may be satisficing for the reasons that Simon (1955) originally proposed, as a rule of thumb that performs adequately across a broad range of environments, rather than finely honing their search strategy to each choice environment they face." This procedural interpretation may suggest that procedural complexity might have an important role in producing satisficing behavior.

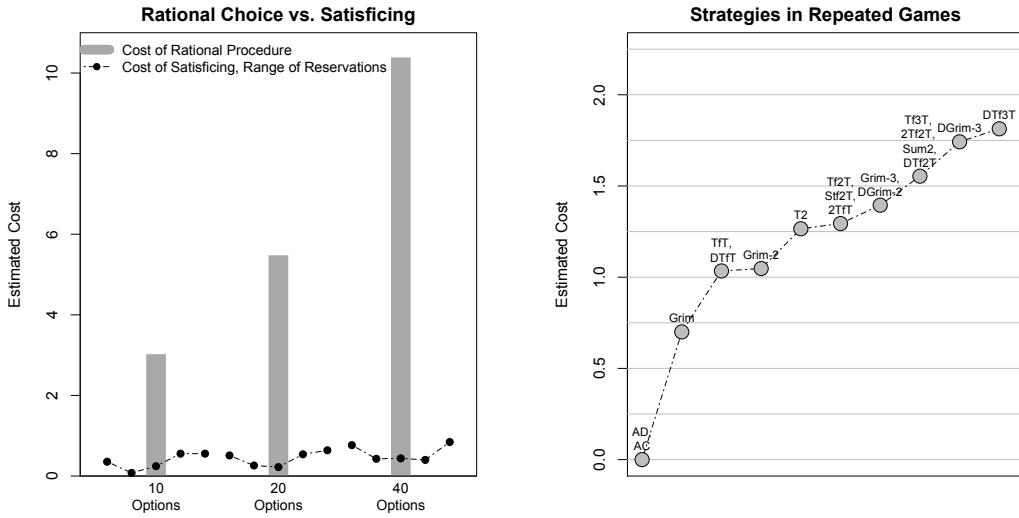


Figure 9. : Two examples of calibrated complexity costs.

Note: The left panel calibrates complexity costs (gray bars) of rationally choosing the highest value item from lists of options of size 10, 20 and 40. For reference, in each case, we plot (dotted lines) the opportunity cost of satisficing (from Caplin, Dean and Martin (2011)) with a range of reservation values (5, 10, 15, 20 and 25). The right hand panel calibrates complexity costs for a set of strategies for repeated prisoner's dilemmas.

is that the effort required to evaluate individual items (in their experiment the number of numbers that must be summed to discern an item's value) has an influence on mistakes and rates of satisficing. A natural way of understanding this in the context of automata models, is that in some settings evaluating *events* can be costly: each time subjects in Caplin, Dean and Martin (2011) evaluate a new item (each time they make a transition, in the language of Salant (2011)), they must perform a computation and as these computations grow costly, the costs of implementing elaborate procedures (like rational choice) must grow. Such costs are set extremely low in our experiment, where subjects immediately (with low required cognitive effort) know what event has occurred and thus how to transition behavior in response. If our rules had required subjects to add 3 or 7 numbers to determine which letter-event had occurred, our cost estimates would have been more comparable to those faced by subjects in Caplin, Dean and Martin (2011), and we would expect subjects to be more strongly averse to rules that tend to require more event-processing and transitions. From this perspective our cost calibrations probably understate the procedural costs faced by subjects in Caplin, Dean and Martin (2011). As Salant (2011) emphasizes, understanding how the “marginal” complexity costs of evaluating each option interacts with the “fixed” costs of the decision procedure itself (for instance the number of states in the procedure) is an important task for future work.

B. Repeated Games Strategies

Repeated games arise across economics, but the theory suffers from a fundamental indeterminacy, described by the folk theorem: with low enough discount rates, a wide variety of outcomes can generally be supported if players are willing to use sufficiently complex punishment strategies. The automata literature began as an attempt to study how aversion to complex strategies might reduce this indeterminacy, making game theory more predictive.

In the right panel of Figure 9 we calibrate the costs of a number of strategies from repeated prisoner's dilemmas.³⁷ Specifically, we calibrate complexity costs using our estimates of the costs of states, transitions and absorption for most (18) of the strategies discussed and estimated in Fudenberg, Rand and Dreber (2012).³⁸ The results reveal a wide range of complexity costs, with the least complex non-trivial strategy (Grim) costing 1/3 as much as the most complex (DTf3T).

Based on these cost estimates alone (i.e. without taking account of the benefits of strategies or important issues of equilibrium and strategic risk), several predictions about the use of these strategies naturally arise. For one thing, *ceteris paribus*, we would expect the “cheapest” strategies – AC (always cooperate), AD (always defect), Grim (the grim trigger strategy), Tft (tit for tat) and D-Tft (tit for tat beginning with defection) – to be particularly common and for our most “expensive” strategies (e.g. DGrim3, DTf3T, 2Tf2T) to be relatively rare. For another, because there are significant costs attached to implementing strategies more complex than the extremely simple (one-shot Nash) strategy, AD, we should expect to see subjects abandoning this simplest strategy exactly when the parameters of the game increase the potential payoffs from avoiding a mutual defection outcome (a moderate payoff outcome that subjects can guarantee for themselves using a maximally simple rule).

Indeed, both of these patterns are typical in the experimental literature. For instance Dal Bó and Fréchette (2018) review the literature and calculate that our estimated five ‘cheapest’ strategies (AC, AD, Grim, Tft and D-Tft) account for the majority of strategies in almost all papers that estimate strategies, and that they account for more than 3/4 of strategies in the majority of studies (in fact AD, Grim and Tft alone account for more than 70% of strategies in the

³⁷Here AD and AC are “always defect” and “always cooperate,” Grim-N is the Grim rule from the introduction with N periods of counterpart defection preceding own defection. Tit-for-tat specifies cooperating and then defecting after counterpart defection and returning to cooperation the next round if the counterpart does not defect again. (M)Tf(N)T is tit-for-tat that moves from cooperation only after (N) rounds of counterpart defection and then remains in defection for (M) rounds before returning if counterpart is cooperative. Rules that begin with a ‘D’ add an initial period of defection to the beginning of the rule. T2 is a strategy that defects twice in response to a defection.

³⁸For instance, the rule “Grim 3” – cooperate until your counterpart defects three times, then defect forever – has four states, three transitions and an absorbing state. Using estimates from specification (1) in Table 5, we estimate a cost by adding our intercept 0.903 (which estimates the baseline cost of a 2-state, 1 transition, transient rule), $2 \times (0.239 - 0.14)$ (the cost of states plus the interaction with an absorbing state) and $2 \times (0.116 + 0.131)$ (transitions plus the interaction with an absorbing state) and subtracting 0.203 (because of the absorbing state) we get a cost of \$1.40.

majority of papers/treatments). Dal Bó and Fréchette (2011) and Dal Bó and Fréchette (2019) systematically vary the cooperation payoffs and the discount rate (which together determine the potential benefits of using more complex strategies than AD) and find subjects are, indeed, much more likely to abandon AD for richer strategies when either of these parameters rise. In some settings (e.g. prisoner's dilemmas with randomly perturbed actions) the benefits of cooperation are difficult to achieve in any other way than to use relatively complex strategies and it is only in these settings that the literature documents any significant use of moderately complex, costly strategies like Grim-3 and Tf3T (rather than their simpler cousins Grim and TfT) (Fudenberg, Rand and Dreber, 2012).³⁹ Under the lens of our findings, it seems that subjects actively avoid complex rules when possible, deploying them only when the gains to doing so are sufficiently high.

Again, we emphasize that further research is needed before we can draw firm conclusions on the role of procedural complexity in strategy selection. This is particularly important in rich settings like repeated games where equilibrium and strategic uncertainty considerations cross with complexity costs in surely complicated ways. New experiments that correlate individual complexity cost measures with strategy choice and especially novel designs built on environments capable of separating the role of procedural complexity from that of other key forces will be important for this task.⁴⁰

V. Conclusion

We use a new type of experiment to identify the algorithmic characteristics of rules that make them complex and therefore costly to implement. We find evidence of significant structure in the determinants of complexity, with a rule's algorithmic characteristics (particularly the number of states required to implement it) significantly influencing its complexity costs. We also find evidence that the way a rule can be represented (for instance whether it can be simplified using counting), familiarity with a rule acquired via learning, and the context in which a rule must be implemented (e.g. physically vs. mentally) have additional influence on complexity costs.

Our methods and iterations on them may allow us to use procedural complexity as a way of explaining, modeling and predicting departures from standard eco-

³⁹Fudenberg, Rand and Dreber (2012) report experimental prisoner's dilemmas run both with and without noisy implementation of actions (i.e. with and without a small probability the computer alters the subject's actions), under identical payoff parameters. They too find relatively low usage of complex strategies without noisy implementation, but find that with noisy implementation a majority of subjects use strategies more complex than TfT. Simpler strategies in this setting can produce costly punishment cycles (based purely on random events) that more complex strategies allow the subject to avoid.

⁴⁰An example of the type of design that might be expanded upon for this task is Jones (2014) who artificially alters prisoner's dilemmas in such a way as to increase the state complexity of any cooperative strategy. Mirroring our findings in support of state complexity, Jones (2014) finds that increasing the state complexity of cooperative strategies reduces the rate of cooperation observed in prisoner's dilemmas. Taking a similar approach in other games and choice problems and relating the results to the direct measurement techniques we introduce here seems like a promising template for future procedural complexity research.

nomics predictions documented in behavioral and institutional economics. Testing the ability of these types of measurements to explain behavioral anomalies, equilibrium selection and policy failures – both through new diagnostic experimental designs and correlational studies – is an important next step in developing this research. Furthermore, careful work examining to what degree richer algorithmic descriptions of rules (e.g. pushdown automata, Turing machines) might improve our ability to predict and explain complexity and its costs seems particularly important. Finally, although we have applied our methods to the complexity costs of rules, in principle they may be useful for measuring the costs of other types of behavior and thus may have a much broader scope of application.

REFERENCES

Abeler, Johannes, and Simon Jäger. 2015. “Complex Tax Incentives.” *American Economic Journal: Economic Policy*, 7(3): 1–28.

Abreu, Dilip, and Ariel Rubinstein. 1988. “The Structure of Nash Equilibrium in Repeated Games with Finite Automata.” *Econometrica*, 56(6): 1259–1281.

Alaoui, Larbi, and Antonio Penta. 2018. “Cost-Benefit Analysis in Reasoning.”

Banks, Jeffrey S, and Rangarajan K Sundaram. 1990. “Repeated Games, Finite Automata, and Complexity.” *Games and Economic Behavior*, 2(2): 97–117.

Bossaerts, Peter, and Carsten Murawski. 2017. “Computational Complexity and Human Decision-Making.” *Trends in Cognitive Sciences*, 21(12): 917–929.

Campbell, Donald J. 1988. “Task Complexity: A Review and Analysis.” *The Academy of Management Review*, 13: 40–52.

Caplin, Andrew. 2016. “Measuring and Modeling Attention.” *Annual Review of Economics*, 8: 379–403.

Caplin, Andrew, Dániel Csaba, John Leahy, and Oded Nov. 2019. “Rational Inattention, Competitive Supply, and Psychometrics.”

Caplin, Andrew, Mark Dean, and Daniel Martin. 2011. “Search and Satisficing.” *American Economic Review*, 101(7): 2899–2922.

Cason, Timothy N, and Charles Plott. 2014. “Misconceptions and Game Form Recognition: Challenges to Theories of Revealed Preferences and Framing.” *Journal of Political Economy*, 122: 1235–1270.

Chatterjee, Kalyan, and Hamid Sabourian. 2000. “Multiperson Bargaining and Strategic Complexity.” *Econometrica*, 68(6): 1491–1509.

Chatterjee, Kalyan, and Hamid Sabourian. 2009. “Game Theory and Strategic Complexity.” *Encyclopedia of Complexity and Systems Science*, 4098–4114.

Chauvin, Kyle P. 2020. “Euclidean Properties of Bayesian Updating.”

Cooper, David J., and John H. Kagel. 2010. “The Role of Context and Team Play in Cross-Game Learning.” *Journal of the European Economic Association*, 7(5): 1101–1139.

Cosmides, Leda. 1987. “The Logic of Social Exchange: Has Natural Selection Shaped How Humans Reason? Studies with the Wason Selection Task.” *Cognition*, 31: 187–276.

Dal Bó, Pedro, and Guillaume R Fréchette. 2011. “The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence.” *American Economic Review*, 101(1): 411–429.

Dal Bó, Pedro, and Guillaume R Fréchette. 2018. “On the Determinants of Cooperation in Infinitely Repeated Games: A Survey.” *Journal of Economic Literature*, 56(1): 60–114.

Dal Bó, Pedro, and Guillaume R Fréchette. 2019. “Strategy Choice In The Infinitely Repeated Prisoners Dilemma.” *American Economic Review*, forthcoming.

Dean, Mark, and Nathaniel Neligh. 2019. “Experimental Tests of Rational Inattention.”

Dewam, Ambuj, and Nathaniel Neligh. 2020. “Estimation Information Cost Functions in Models of Rational Inattention.” *Journal of Economic Theory*, 187.

Echenique, Federico, Daniel Golovin, and Adam Wierman. 2011. “Complexity and Economics: Computational Constraints May Not Matter Empirically.” *ACM SIGecom Exchanges*, 10(1): 2–5.

Ergin, Haluk, and Todd Sarver. 2010. “A Unique Costly Contemplation Representation.” *Econometrica*, 78(4): 1285–1339.

Fershtman, Chaim, and Ehud Kalai. 1993. “Complexity Considerations and Market Behavior.” *The RAND Journal of Economics*, 24(2): 224–235.

Fudenberg, Drew, David Rand, and Anna Dreber. 2012. “Slow to Anger and Fast to Forgive: Cooperation in an Uncertain World.” *American Economic Review*, 102(2): 720–749.

Gale, Douglas, and Hamid Sabourian. 2005. “Complexity and Competition.” *Econometrica*, 73(3): 739–769.

Gill, David, and Victoria L Prowse. 2018. “Strategic Complexity and the Value of Thinking.”

Greiner, Ben. 2015. “Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE.” *Journal of the Economic Science Association*, 1(1): 114–125.

Halevy, Yoram, and Guy Mayraz. 2020. “Modes of Rationality: Act versus Rule-Based Decisions.”

Hopcroft, John E., and Jeffrey D. Ullman. 1979. *Introduction to Automata Theory: Languages and Computation*. Addison-Wesley.

Jones, Matthew T. 2014. “Strategic Complexity and Cooperation: An Experimental Study.” *Journal of Economic Behavior & Organization*, 106: 352–366.

Kalai, Ehud. 1990. “Bounded Rationality and Strategic Complexity in Repeated Games.” *Game Theory and Applications*, 131–157.

Kalai, Ehud, and Eilon Solan. 2003. “Randomization and Simplification in Dynamic Decision-Making.” *Journal of Economic Theory*, 111(2): 251–264.

Kalai, Ehud, and William Stanford. 1988. “Finite Rationality and Interpersonal Complexity in Repeated Games.” *Econometrica*, 56(2): 397–410.

Kool, Wouter, and Matthew Botvinick. 2018. “Mental Labour.” *Nature Human Behavior*, 2: 899–908.

Lipman, Barton L, and Sanjay Srivastava. 1990. “Informational Requirements and Strategic Complexity in Repeated Games.” *Games and Economic Behavior*, 2(3): 273–290.

Liu, Peng, and Zhizhong Li. 2012. “Task Complexity: A Review and Conceptualization Framework.” *International Journal of Industrial Ergonomics*, 42(6): 553–568.

Murawski, Carsten, and Peter Bossaerts. 2016. “How Humans Solve Complex Problems: The Case of the Knapsack Problem.” *Scientific reports*, 6: 34851.

Neyman, Abraham. 1985. “Bounded Complexity Justifies Cooperation in the Finitely Repeated Prisoners’ Dilemma.” *Economics Letters*, 19(3): 227–229.

Nielsen, Kirby, and John Rehbeck. 2020. “When Choices are Mistakes.”

Ortoleva, Pietro. 2013. “The Price of Flexibility: Towards A Theory of Thinking Aversion.” *Journal of Economic Theory*, 148(3): 903–934.

Rubinstein, Ariel. 1986. “Finite Automata Play the Repeated Prisoner’s Dilemma.” *Journal of Economic Theory*, 39(1): 83–96.

Sabourian, Hamid. 2004. “Bargaining and Markets: Complexity and the Competitive Outcome.” *Journal of Economic Theory*, 116(2): 189–228.

Salant, Yuval. 2011. “Procedural Analysis of Choice Rules with Applications to Bounded Rationality.” *American Economic Review*, 101(2): 724–748.

Shenhav, Amitai, Sebastian Musslick, Falk Lieder, Wouter Kool, Thomas L. Griffiths, Jonathan D. Cohen, and Matthew M. Botvinick. 2017. “Toward a Rational and Mechanistic Account of Mental Effort.” *Annual Review of Neuroscience*, 40: 99–124.

Sherstyuk, Katerina. 2011. “Complexity and Bidder Behavior in Iterative Auctions.” *Economics Bulletin*, 31(4): 2769–2776.

Simon, Herbert A. 1955. “A Behavioral Model of Rational Choice.” *The Quarterly Journal of Economics*, 69(1): 99–118.

Sims, Christopher A. 2003. “Implications of Rational Inattention.” *Journal of Monetary Economics*, 50(3): 665–690.

Wilson, Andrea. 2014. “Bounded Memory and Biases in Information Processing.” *Econometrica*, 82(6): 2257–2294.