A Peculiar Connection Between the Axiom of Choice and Predicting the Future

Christopher S. Hardin and Alan D. Taylor

1. INTRODUCTION. We often model systems that change over time as functions from the real numbers $\mathbb{R}$ (or a subinterval of $\mathbb{R}$) into some set S of states, and it is often our goal to predict the behavior of these systems. Generally, this requires rules governing their behavior, such as a set of differential equations or the assumption that the system (as a function) is analytic. With no such assumptions, the system could be an arbitrary function, and the values of arbitrary functions are notoriously hard to predict. After all, if someone proposed a strategy for predicting the values of an arbitrary function based on its past values, a reasonable response might be, "That is impossible. Given any strategy for predicting the values of an arbitrary function, one could just define a function that diagonalizes against it: whatever the strategy predicts, define the function to be something else." This argument, however, makes an appeal to induction: to diagonalize against the proposed strategy at a point t, we must have already defined our function for all $s < t$ in order to determine what the strategy would predict at t. So, the argument would only be valid if $\mathbb{R}$ were well-ordered, but $\mathbb{R}$ is emphatically not well-ordered.

In fact, the lack of well-orderedness in the reals can be exploited to produce a very counterintuitive result: there *is* a strategy for predicting the values of an arbitrary function, based on its previous values, that is almost always correct. Specifically, given the values of a function on an interval $(-\infty, t)$, the strategy produces a guess for the values of the function on $[t, \infty)$, and at all but countably many t, there is an $\varepsilon > 0$ such that the prediction is valid on $[t, t + \varepsilon)$. Noting that any countable set of reals has measure 0, we can restate this informally: at almost every instant t, the strategy predicts some "ε-glimpse" of the future.

We should emphasize that these results do not give a practical means of predicting the future, just as the time dilation one would experience standing near the event horizon of a black hole does not give a practical time machine. Nevertheless, we choose this presentation because we find it the most interesting, as well as pedagogically useful. For instance, "predicting the present" is a very natural way to think of the problem of guessing the value of $f(t)$ based on $f|(-\infty, t)$.

2. THE μ-STRATEGY. Fix sets T and S, with $|S| \geq 2$, and a binary relation $\lhd$ on T. (The situation of the introduction will be the special case where $(T, \lhd) = (\mathbb{R}, <)$.)

Definition 2.1. Let ${}^T S$ denote the set of all functions from T to S. We call the elements of ${}^T S$ *scenarios*. For each $t \in T$, we define the equivalence relation $\approx_t$ on ${}^T S$ by $f \approx_t g$ iff f and g agree on $\lhd t = \{ s \in T \mid s \lhd t \}$, the set of $\lhd$-predecessors of t. We let $[f]_t$ denote the equivalence class of f under $\approx_t$.

Before proceeding, we make two related interpretations of the above definition: one in terms of hats, and the other in terms of time. In the hat interpretation, T represents some set of agents, each of whom wears a hat whose color is an element of S. A scenario specifies the colors of the hats. For agents s and t, $s \lhd t$ means that agent

t can see the color of agent s's hat, and $f \approx_t g$ means that scenarios f and g are indistinguishable to agent t, in the sense that f and g assign the same colors to the hats that t can see. Given a scenario f for the coloring of the hats, $[f]_t$ is the set of scenarios consistent with what t can see. Under this interpretation, we will examine how well the agents can guess the colors of hats they cannot see (in particular, the colors of their own hats).

In the time interpretation, T represents instants of time, ordered by $\lhd$. A scenario describes the evolution of a system over time, and $f \approx_t g$ means that scenarios f and g are the same up to (but not necessarily at) instant t. We will also treat T as a set of agents, in this case by identifying each instant t with an agent who can see the past at that instant. Under this interpretation, we will examine how well the agents can predict the present or even the future while knowing only the past.

We will often fix an element $v \in {}^T S$ which we consider to be the *true scenario*. Within this scenario, each agent t will attempt to guess v while only being able to see the values of v on $\lhd t$. This amounts to guessing v from $[v]_t$, since $[v]_t$ will be the set of scenarios that agent t considers possible. One strategy for guessing is as follows.

Definition 2.2. Let $\mathcal{O} = \{ [v]_t \mid v \in {}^T S \text{ and } t \in T \}$. A *strategy* is a function $g : \mathcal{O} \to {}^T S$ such that $g([v]_t) \in [v]_t$ for any $v \in {}^T S$ and $t \in T$. Fix a well-ordering $\preceq$ of ${}^T S$ (and note that this requires the Axiom of Choice). The *μ-strategy* is the strategy $\mu : \mathcal{O} \to {}^T S$ defined by letting $\mu([f]_t)$ be the $\preceq$-least element of $[f]_t$. We abbreviate $\mu([v]_t)$ by $\langle v \rangle_t$.

Given a true scenario $v \in {}^T S$, an agent t plays the μ-strategy by guessing that the true scenario is $\langle v \rangle_t$. If we interpret $f \prec g$ as meaning that f is "simpler" than g in some sense, then the μ-strategy can be thought of as a formalization of Occam's razor: each agent chooses the "simplest" scenario that is consistent with what the agent can observe. (Of course, if $f \prec g$ means that f is more complicated than g, then this is the opposite of Occam's razor.)

How good is the μ-strategy? We consider this question in the following sections, first in terms of predicting the present from the past, and then in terms of predicting the present and some of the future from the past.

3. PREDICTING THE PRESENT.

Fix a true scenario $v \in {}^T S$, and let

$$W_0 = \{ t \in T \mid \langle v \rangle_t(t) \neq v(t) \}.$$

In terms of hats, this is the set of agents who guess their own hat color incorrectly when using the μ-strategy; in terms of time, this is the set of instants at which the μ-strategy incorrectly predicts the present from the past.

Theorem 3.1. *If $\lhd$ is transitive, then W_0 is well-founded in $\lhd$; that is, W_0 has no infinite descending $\lhd$-chain.*

Proof. Suppose there is a sequence of agents $t_0, t_1, \ldots$ with $t_i \in W_0$ and $\cdots \lhd t_2 \lhd t_1 \lhd t_0$. For any i, since $t_{i+1} \lhd t_i$ and $\langle v \rangle_{t_i} \approx_{t_i} v$, we have

$$\langle v \rangle_{t_i}(t_{i+1}) = v(t_{i+1}) \neq \langle v \rangle_{t_{i+1}}(t_{i+1}),$$

so $\langle v \rangle_{t_i} \neq \langle v \rangle_{t_{i+1}}$. Furthermore, since $\lhd$ is transitive, we have $\lhd t_{i+1} \subseteq \lhd t_i$, so $[v]_{t_i} \subseteq [v]_{t_{i+1}}$ and $\langle v \rangle_{t_{i+1}} \preceq \langle v \rangle_{t_i}$.

We now have $\langle v \rangle_{t_0} \succ \langle v \rangle_{t_1} \succ \langle v \rangle_{t_2} \succ \cdots$, contradicting the fact that $\preceq$ well-orders ${}^T S$. ■

Corollary 3.2. *If $(T, \lhd)$ is a strict linear order, then W_0 is a well-ordered subset of T.*

Corollary 3.3. *If $(T, \lhd)$ is a strict linear order with no infinite increasing chains, then W_0 is finite.*

Proof. An infinite linear order must have an infinite increasing chain or an infinite decreasing chain; W_0 can have neither. ■

Corollary 3.4. *If $T = \mathbb{R}$ and $\lhd$ is $<$, then W_0 is countable, has measure* 0*, and is nowhere dense.*

Proof. Since W_0 is a well-ordered subset of $\mathbb{R}$, there will always be a rational number between any element $t \in W_0$ and the next element of W_0 (unless t is the last element of W_0, which can happen at most once). So W_0 is countable, and consequently has measure 0.

Any well-ordered subset of $\mathbb{R}$ is nowhere dense, because its closure is again a well-ordered set, and the interior of a well-ordered set must be empty, since nonempty open sets have no least element. ■

What Corollary 3.4 tells us is that, if we model the universe as a function from the real numbers into some set of states, then the μ-strategy will correctly predict the present from the past on a set of full measure. (In the following section, we show that, on a set of full measure, it correctly predicts some of the future as well.) Note that these results concerning $T = \mathbb{R}$ are also valid when T is any interval of reals.

One needs to be cautious about interpreting this as meaning that the μ-strategy is correct with probability 1. For a fixed true scenario, if one randomly selects an instant t in the interval [0,1] (or in $\mathbb{R}$, under a suitable probability distribution), then Corollary 3.4 does tell us that the μ-strategy will be correct at t with probability 1. However, if one fixes the instant t, and randomly selects a true scenario, then the probability that the μ-strategy is correct at t under that scenario might be 0 or might not even exist, depending on how one defines the notion of a random scenario.

The following sharpness result shows that when $\lhd$ is transitive, Theorem 3.1 is the best one can do.

Theorem 3.5. *Suppose $\lhd$ is transitive, $g : \mathcal{O} \to {}^T S$ is a strategy, and $W \subseteq T$ is well-founded in $\lhd$. Then there is a scenario v in which strategy g is wrong at every point in W; that is, $W \subseteq \{ t \in T \mid g([v]_t)(t) \neq v(t) \}$.*

Proof. First, define v arbitrarily on the complement of W. For the points in W, we may define v by induction, since W is well-founded. For $t \in W$, if we have already defined v on the $\lhd$-predecessors of t, then $[v]_t$ is determined, and we define $v(t)$ so that $v(t) \neq g([v]_t)(t)$. ■

4. PREDICTING THE FUTURE. With a more careful analysis, we can show that the μ-strategy often succeeds at predicting some of the future. Let v and W_0 be as in Section 3, and let

$$W_1 = \{ t \in T \mid \langle v \rangle_t \neq \langle v \rangle_{t'} \text{ whenever } t \lhd t' \},$$

the set of instants where the μ-strategy's guess differs from all later guesses. Note that when $\lhd$ is transitive, the μ-strategy never reverts to an older guess after a change is made; in other words, if $t \lhd t' \lhd t''$ and $\langle v \rangle_t \neq \langle v \rangle_{t'}$, then $\langle v \rangle_t \neq \langle v \rangle_{t''}$, since $\langle v \rangle_t \prec \langle v \rangle_{t'} \preceq \langle v \rangle_{t''}$.

Proposition 4.1. $W_0 \subseteq W_1$.

Proof. Suppose $t \in W_0$, and thus $\langle v \rangle_t(t) \neq v(t)$. For any $t' \in T$ with $t \lhd t'$, we have $\langle v \rangle_{t'}(t) = v(t) \neq \langle v \rangle_t(t)$, so $\langle v \rangle_t \neq \langle v \rangle_{t'}$. Therefore $t \in W_1$. ■

Theorem 4.2. *If $\lhd$ is transitive, then W_1 is well-founded in $\lhd$.*

Proof. The proof of Theorem 3.1 applies, except that $\langle v \rangle_{t_i} \neq \langle v \rangle_{t_{i+1}}$ is now immediate from the definition of W_1. ■

Note that, because of the similarity of Theorem 4.2 to Theorem 3.1, Corollaries 3.2–3.4 also apply to W_1.

We again turn our attention to the case $(T, \lhd) = (\mathbb{R}, <)$.

Theorem 4.3. *For $t \in \mathbb{R}$, say that the μ-strategy* guesses well at t *if there exists an $\varepsilon > 0$ such that $\langle v \rangle_t$ and v agree on $[t, t+\varepsilon)$. Then the set of $t \in \mathbb{R}$ where the μ-strategy guesses well has full measure.*

The intuition behind the following proof is that, in any interval disjoint from W_1, the μ-strategy does not change its guess—that is, the map $t \mapsto \langle v \rangle_t$ is constant—but it would if it were ever incorrect about the value of v within the interval; so the one guess made by the μ-strategy within the interval must agree with v until the end of the interval.

Proof. Take any $t \in \mathbb{R} \setminus W_1$. By the definition of W_1, there must be some $t' > t$ such that $\langle v \rangle_t = \langle v \rangle_{t'}$. For any $u < t'$, we have

$$\langle v \rangle_t(u) = \langle v \rangle_{t'}(u) = v(u),$$

so $\langle v \rangle_t$ and v agree on $(-\infty, t')$; in particular, $\langle v \rangle_t$ and v agree on $[t, t+\varepsilon)$ where $\varepsilon = t' - t$. So the μ-strategy guesses well on $\mathbb{R} \setminus W_1$, which has full measure by Corollary 3.4 (with W_1 in place of W_0). ■

5. A REFINEMENT. Again, let $T = \mathbb{R}$ and fix a true scenario $v \in {}^{\mathbb{R}}S$, but suppose that the agents are only able to see the recent past. Specifically, we redefine $\approx_t$ (for this section only) by

$$f \approx_t g \iff (\exists\, s < t)\ f \text{ and } g \text{ agree on } (s, t).$$

For any $s < t$, if we take any given scenario and modify it on the interval $(-\infty, s)$, t will not be able to distinguish the new scenario from the original. So, informally, t sees only an infinitesimal portion of the past.

Although $\approx_t$ is not induced by any relation $\lhd$, we can still make sense of $[v]_t$, the μ-strategy, and $\langle v \rangle_t$, since these were defined in terms of $\approx_t$. And again, we say the μ-strategy *guesses well at t* if there is an interval $[t, t+\varepsilon)$, $\varepsilon > 0$, such that $\langle v \rangle_t$ and v agree on $[t, t+\varepsilon)$.

Let

$$W = \{\, t \in \mathbb{R} \mid \text{the } \mu\text{-strategy does not guess well at } t \,\}.$$

Theorem 5.1. *The set W is countable, has measure* 0*, and is nowhere dense.*

So, even if the agents can see only the recent past, the μ-strategy still almost always correctly predicts some of the future. Daniel Velleman contributed a major simplification to the following proof.

Proof. We first show W is countable. For each rational number q, let

$$W_q = \{\, t \in W \mid q < t \text{ and } \langle v \rangle_t \text{ and } v \text{ agree on } (q, t) \,\}.$$

For each $t \in W$, $\langle v \rangle_t$ and v must agree on (s, t) for some $s < t$; letting $q \in (s, t)$ be rational, we have $t \in W_q$. So W is the union of the sets W_q.

Suppose $t, t' \in W_q$ with $t < t'$. Because $t \in W$, there must be some $u \in [t, t')$ such that $\langle v \rangle_t(u) \neq v(u)$; we then have $\langle v \rangle_t(u) \neq v(u) = \langle v \rangle_{t'}(u)$, so $\langle v \rangle_t \neq \langle v \rangle_{t'}$. Also, because $\langle v \rangle_{t'}$ and v agree on (q, t'), and $q < t < t'$, we have $\langle v \rangle_{t'} \approx_t v$, so $\langle v \rangle_t \preceq \langle v \rangle_{t'}$. It follows, as in the proof of Theorem 3.1, that W_q is well-ordered and hence countable.

Being the union of countably many countable sets, W is countable and has measure 0.

To show W is nowhere dense, suppose, by way of contradiction, that W is dense on an interval (a, b) with $a < b$. Then there exists $t_0 \in W \cap (a, b)$, and $s_0 < t_0$ such that $\langle v \rangle_{t_0}$ and v agree on (s_0, t_0); we assume without loss of generality that $a \leq s_0$. Continuing in this manner, we define sequences s_i and t_i in $\mathbb{R}$ such that $t_i \in W$, $\langle v \rangle_{t_i}$ and v agree on (s_i, t_i), and $s_i \leq s_{i+1} < t_{i+1} < t_i$. By the same argument as above, we have $\langle v \rangle_{t_{i+1}} \neq \langle v \rangle_{t_i}$ and $\langle v \rangle_{t_{i+1}} \preceq \langle v \rangle_{t_i}$, giving us an infinite descending chain $\langle v \rangle_{t_0} \succ \langle v \rangle_{t_1} \succ \cdots$, contradicting the fact that $\preceq$ is a well-ordering. ∎

6. THE ROLE OF AC. Although we used the Axiom of Choice (AC) when defining the μ-strategy, one might wonder if a different approach could yield similar results without using AC. In nontrivial cases, the answer is no.

Theorem 3.1 is interesting only when T has an infinite descending chain and $|S| \geq 2$. In such cases, if there is any strategy for which the analogue of Theorem 3.1 holds, then there exists a set of reals that does not have the property of Baire (and is not Lebesgue measurable). We omit the argument here, but it can be carried out in ZF. Assuming the consistency of ZF, it is consistent with ZF that every set of reals has the property of Baire [**1**], so it is consistent with ZF that, in nontrivial cases, there is no strategy for which the analogue of Theorem 3.1 holds. So, some use of AC is necessary to prove the existence of such strategies, in the sense that ZFC is adequate while ZF is not.

7. HAT PUZZLES. There is a well-known puzzle in which finitely many people are placed in a line, a black or white hat is placed on each one's head, and each person can see only the hats of the people toward the back of the line; starting at the front of the line, they must attempt to guess the colors of their own hats, after hearing the previous guesses. (The solution allows all but possibly the first to guess correctly, by using parity.) Yuval Gabay and Michael O'Connor produced several infinitary variations of this puzzle, one of which is the following.

Puzzle. A prison warden lines up infinitely many prisoners, numbered $0, 1, 2, \dots$, in a large prison yard, and places a black or white hat on each prisoner's head. Each prisoner can see only the hats of the higher-numbered prisoners. Each prisoner must venture a guess about his own hat color, without hearing the guesses of any other prisoners (though they may collaborate on a strategy before the hats are placed). Those that guess correctly are released. How well can the prisoners do at guessing? (Specifically, is there a strategy that guarantees at least one prisoner will be correct? Is there a strategy that guarantees at most finitely many are wrong?)

They solved the puzzle by a different method than we use below, and left open a further variation: What if the prisoners are numbered by ordinals in ω_1? That puzzle led to the work here.

Our solution is that, if the ordering of the prisoners is any ordinal, then the prisoners can guarantee that all but finitely many guess correctly, by applying the μ-strategy. Specifically, for the variation where the prisoners are numbered by elements of an ordinal α, we let $T = \alpha$, and say $x \lhd y \iff y < x$; Corollary 3.3 then implies that all but finitely many prisoners will guess correctly when the μ-strategy is used.

ACKNOWLEDGMENTS. The authors gratefully acknowledge Gabay and O'Connor for these puzzles, and two anonymous referees for several suggestions that have been incorporated.

REFERENCES

1. H. Judah and S. Shelah, Baire property and axiom of choice, *Israel J. Math.* **84** (1993) 435–450.

CHRISTOPHER S. HARDIN received his B.A. from Amherst College in 1998 and his Ph.D. from Cornell University in 2005. He is currently a visiting professor at Smith College.
Department of Mathematics and Statistics, Smith College, Northampton, MA 01063
chardin@smith.edu

ALAN D. TAYLOR received his B.A. from the University of Maine in 1969 and his Ph.D. from Dartmouth College in 1975. He is Marie Louise Bailey Professor of Mathematics at Union College.
Department of Mathematics, Union College, Schenectady, NY 12308
taylora@union.edu