

INDIVIDUAL LEARNING ABOUT CONSUMPTION

TODD W. ALLEN

PriceWaterhouseCoopers

CHRISTOPHER D. CARROLL

NBER

and

The Johns Hopkins University

The standard approach to modeling consumption/saving problems is to assume that the decisionmaker is solving a dynamic stochastic optimization problem. However, under realistic descriptions of utility and uncertainty, the optimal consumption/saving decision is so difficult that only recently have economists managed to find solutions, using numerical methods that require previously infeasible amounts of computation. Yet, empirical evidence suggests that household behavior conforms fairly well with the prescriptions of the optimal solution, raising the question of how average households can solve problems that economists, until recently, could not. This paper examines whether consumers might be able to find a reasonably good rule-of-thumb approximation to optimal behavior by trial-and-error methods, as Milton Friedman proposed long ago. We find that such individual learning methods can reliably identify reasonably good rules of thumb only if the consumer is able to spend absurdly large amounts of time searching for a good rule.

Keywords: Learning, Consumption, Saving, Uncertainty, Buffer-Stock Saving

1. INTRODUCTION

The last decade has seen an explosion of research on learning and evolutionary dynamics in the game theory literature [for recent surveys see, e.g., Gale (1996), Samuelson (1997), or Young (1997)], and the development of substantial literatures on learning and strategy evolution in macroeconomic models [for surveys, see Sargent (1993), Marimon (1997), or Evans and Honkapohja (1999)] and finance [Arthur et al. (1997), Lettau (1997)]. However, there has been remarkably little work on the role of learning in the realm of intertemporal choice problems such as consumption/saving and investment decisions.¹

The traditional approach to modeling intertemporal decisions has been to assume that economic agents are solving a mathematical dynamic programming problem. Long ago, Milton Friedman (1953) defended the optimization assumption by

We would like to thank the editors of this volume, along with John Miller, Scott Page, and seminar participants at the Santa Fe Institute for helpful comments on earlier drafts of this paper. Address correspondence to: Christopher Carroll, Department of Economics, Johns Hopkins University, Baltimore, MD 21218-2685, USA; e-mail: ccarroll@jhu.edu.

arguing that agents could learn roughly optimal behavior by a process of trial and error. Yet nearly 50 years later, in the domain of intertemporal choice, Friedman's proposition remains a largely unexamined assertion rather than a conclusion based on either empirical evidence or models of learning.

The principal reason that there has been little examination of Friedman's learning hypothesis in the context of intertemporal problems is probably that such problems are astonishingly difficult to solve. Only in the past 10 years or so, starting with the work of Zeldes (1989), have economists finally managed (using numerical methods requiring previously infeasible amounts of computer time) to solve the optimal consumption problem under realistic specifications of uncertainty and plausible assumptions about the utility function. Solving and understanding these models, and discovering that their implications fit the data surprisingly well, has occupied the minds and time of consumption researchers for much of the past decade.

This paper begins an investigation of whether consumers who do not understand dynamic stochastic optimization theory and do not have access to very fast computers might still be able to learn roughly optimal behavior by trial and error, as Friedman argued so long ago.

The first contribution of the paper is to show that, although finding the exactly correct nonlinear consumption policy rule (as economists have done) is an extraordinarily difficult mathematical problem, the exactly correct rule can be very closely approximated (in utility terms) by a linear form that seems simple enough that consumers could plausibly learn it by trial and error, because both the slope and the intercept have highly intuitive interpretations (the intercept determines the target level of wealth, and the slope determines "how hard" the consumer tries to get back to his target wealth when away from it). The conceptual and mathematical simplicity of the linear approximate solution to the buffer-stock saving problem makes it a natural framework to use as a proving ground for models of learning about intertemporal choice under uncertainty, just as the deceptively simple Prisoner's Dilemma problem is the prototype for models of learning in a game theory context (or, ranging further, as *E. coli* is a simple proving ground for biological research).

The second contribution of the paper, however, is to show that even the simplified linear consumption function is enormously difficult to find by trial and error. The difficulty stems from our assumption that the consumer cannot directly perceive the value function associated with a given consumption rule, but instead must evaluate the consumption rule by living with it for long enough to get a good idea of its performance. In the consumption problem, the decision to consume a bit more this year implies lower wealth next year, and the year after, and so on into the distant future. Because the consequences of today's actions are spread out over a very long time horizon, it is necessary to experience a long time horizon in order to reliably determine the value of any candidate consumption rule. Furthermore, empirically plausible amounts of uncertainty make the problem much more difficult because the effects of a given *shock* are spread out over time as well. Thus, it takes a very large amount of experience with each potential consumption rule to get an accurate sense of how good or bad that rule is. This situation is a strong contrast with most

of the existing literature on learning in both macroeconomics and game theory, where the typical assumption is that all of the consequences of a choice made at time t are observed immediately. Intertemporal problems are evidently orders of magnitude more difficult.

Despite the extraordinary difficulty of finding a reasonably good consumption rule, empirical evidence suggests that typical households do engage in buffer-stock saving behavior.² The question of how consumers come by their consumption rules therefore remains. Perhaps the most plausible answer involves “social learning”: Rather than relying solely on their own (insufficient) experience, people observe the experiences of others and can learn from such observation and direct social communication. However, the existing literature on social learning [for surveys, see Bikhchandani et al. (1998) or Gale (1996)] suggests that social learning mechanisms are by no means guaranteed to converge on the optimum. Exploring the circumstances under which social learning processes do and do not lead a population to converge on reasonably optimal behavior promises to be an interesting task for future work.

2. BACKGROUND AND LITERATURE SUMMARY

Because the optimal consumption/saving problem does not have an analytical solution under plausible specifications of utility and uncertainty, until very recently economists usually solved versions of the model in which consumers either had unrealistic (quadratic) preferences for which uncertainty does not affect consumption, or had plausible [constant relative risk aversion (CRRA)] preferences but faced no uncertainty.

This certainty equivalent (CEQ) model has been tested exhaustively. An influential summary of the literature [Deaton (1992)] has suggested that the model fails in at least three ways. First, a large literature dating from the 1950's and 1960's and extending through Hall and Mishkin (1982), McCarthy (1995), Parker (1999), and Souleles (1999) has consistently estimated a marginal propensity to consume greater than 0.2; since the CEQ model generally implies MPC's of less than 0.05, these results have been interpreted as suggesting the presence of some consumers who always consume their entire income, either because they are rational but liquidity constrained or because they are simply irrational. Second, another large literature has tested the CEQ model's prediction that the marginal propensity to consume out of human wealth is the same as the MPC out of current wealth, and consistently has found consumption and saving to be largely unresponsive to information about future income.³ Third, a vast literature estimating Euler equations arose from Hall (1978). A recent survey article in the *Journal of Economic Literature* by Browning and Lusardi (1996) summarized over 25 studies using microeconomic data to estimate an Euler equation derived from standard versions of the model. Most of the studies rejected the Euler equation, usually in favor of a model in which some consumers simply blindly set consumption equal to income. A final failure of the CEQ model is that it provides no explanation for one of the

central and robust findings from household wealth surveys: All such surveys, from the early 1960's to the most recent (1998) triennial *Survey of Consumer Finances*, have found that the median household at every age before about 50 typically holds total nonhousing net assets worth somewhere between a few weeks' worth and a few months' worth of income [Carroll (1997a)].

Ironically, when advances in computer technology finally permitted numerical solution of the optimal consumption problem under realistic assumptions about uncertainty and preferences,⁴ all of these supposed rejections of rationality turned out to be *consistent* with dynamic optimization after all! Under some plausible combinations of parameter values, optimal behavior is for consumers to aim to hold a target buffer stock of liquid assets equivalent to a few weeks or months' worth of consumption, and once the target wealth is achieved to set consumption on average equal to average income. Even with a time preference rate as low as 0.04, the marginal propensity to consume out of transitory income can be 40% or higher, the propensity to consume out of human wealth can be close to zero, and standard Euler equation tests of consumption behavior "fail" in ways that can replicate the whole range of empirical failures of the Euler equation. [See Carroll (1992, 1997a,b) for details.] Uncertainty and the consequent precautionary saving motive thus turn out to modify optimal behavior profoundly from what was taken by economists to embody "rationality" from the 1950's through the late 1980's.

In a way, the recent findings can be interpreted as a potential vindication of Friedman's argument that people can grasp the solution to difficult mathematical problems even without mathematical training. Perhaps the embarrassment is that economists for so long failed to see what consumers apparently implicitly know: that buffer-stock saving behavior works reasonably well. However, these findings also raise rather urgently the question of how ordinary consumers appear to solve, even approximately, problems whose solution even now, and even in versions much simpler than the actual problems people face, continue to strain the capabilities of powerful modern computers.⁵ One possible answer is that people may have a powerful inbuilt intuition about the solution to dynamic optimization problems. However, this explanation founders on the observation that economists are people too. If anything, inbuilt mathematical intuitions ought to be stronger for economists than for average consumers because economists are much better mathematicians; yet, economists did not discover the optimality of buffer-stock behavior until fast computers made it possible to solve the problem numerically. Friedman's learning hypothesis seems to be the natural alternative explanation.

3. BUFFER-STOCK SAVING: AN APPROXIMATION

One of the attractive features of the buffer-stock theory of saving is that optimal behavior can be articulated in very simple and intuitive terms: Consumers have a target level for a buffer stock of liquid assets that they use to smooth consumption in the face of an uncertain income stream. If their buffer stock falls below its target, they will consume less than their expected income and liquid assets will

rise, whereas if they have assets in excess of their target they will spend freely and assets will fall.

Despite its heuristic simplicity, the *exact* mathematical specification of optimal behavior is given by a thoroughly nonlinear consumption rule for which there is no analytical formula. Although certain analytical characteristics of the rule can be proven,⁶ it is hard to see how a consumer without a supercomputer and a doctorate could be expected to determine the exact shape of the nonlinear and nonanalytical decision rule.

3.1. An Approximation

Fortunately for consumers, it turns out not to matter much whether they get the fine details of the rule right: Simple and intuitive approximations to the optimal rule can generate utility streams that are only trivially smaller than the utility yielded by the exact and fully nonlinear solution.

For example, consider a consumer in the following circumstances. Utility is derived entirely from consumption and is CRRA, $u(c) = c^{1-\rho}/(1-\rho)$, with $\rho = 3$. Income Y is stochastic, with a 3-point distribution (0.7, 1, 1.3) with probabilities (0.2, 0.6, 0.2), a process chosen to match (very roughly) empirical evidence on the amount of transitory variation in annual household income observed in the *Panel Study of Income Dynamics* [see, e.g., Carroll (1992)]. The consumer cannot borrow, but can save at an interest rate of zero. Finally, the consumer geometrically discounts future utility at the rate $\beta = 0.95$. The traditional approach to modeling consumer behavior is to suppose that the consumer solves the problem

$$\max_{\{C_s\}_t^\infty} E_t \left[\sum_{s=t}^\infty \beta^{s-t} u(C_s) \right],$$

such that

$$\begin{aligned} X_{s+1} &= X_s - C_s + Y_{s+1}, \\ C_s &\leq X_s \quad \forall s, \end{aligned}$$

where X_s is total resources available for consumption [henceforth, following Deaton (1991), “cash-on-hand”]. Of course, as is well known, the objective in this problem can be rewritten in the recursive form

$$V(X_t) = \max_{\{C_s\}} u(C_s) + \beta E_t[V(\tilde{X}_{t+1})], \tag{1}$$

where $V(X_t)$ is the value function reflecting the expected discounted utility that will result if the consumer behaves optimally now and in all future periods.

As noted above, one interesting feature of the solution to this problem is that there will exist a target level of cash on hand \bar{X}^* . Formally, Carroll (1996) shows that if the parameters of the problem satisfy a certain “impatience” condition,⁷

then an \bar{X}^* will exist such that if $X_t > \bar{X}^*$, then $E_t[X_{t+1}] < X_t$, and if $X_t < \bar{X}^*$, then $E_t[X_{t+1}] > X_t$. Assuming $\bar{X}^* \geq 1$, for some f the optimal consumption rule can be rewritten, without loss of generality, as

$$C^*(X) = 1 + f(X - \bar{X}^*). \tag{2}$$

Using the fact that $E_t[\tilde{Y}_{t+1}] = 1$, we know that $E_t[X_{t+1}] = X_t - C_t + 1$. However, at the point at which $X = \bar{X}^*$, we have $E_t[X_{t+1}] = X_t$, which implies that $X_t - C_t + 1 = X_t$, which implies that $C_t = 1$. Hence we know that $f(0) = 0$. Calling $\gamma^* = f'(0)$, a first-order Taylor expansion of equation (2) around the point $X = \bar{X}^*$ is therefore

$$C^*(X) \approx 1 + \gamma^*(X - \bar{X}^*). \tag{3}$$

Define a variable $\theta = \{\gamma, \bar{X}\}$ and define a function

$$C^\theta(X) = \begin{cases} 1 + \gamma(X - \bar{X}) & \text{if } 1 + \gamma(X - \bar{X}) \leq X, \\ X & \text{if } 1 + \gamma(X - \bar{X}) > X, \end{cases}$$

where the second case implements the liquidity constraint.⁸ Now, choose the values of γ and \bar{X} that correspond to the Taylor approximation to equation (2), $\theta^* = \{\gamma^*, \bar{X}^*\}$, yielding the rule $C^{\theta^*}(X)$. The attraction of this rule, in comparison with the exact nonlinear solution $C^*(X)$, is that it produces a complete plan of behavior that is characterized by only two parameter values, \bar{X}^* and γ^* . Furthermore, it is an approximation that will by construction be close to the true consumption rule in the neighborhood of the target level of wealth; if actual wealth tends to stay relatively close to target wealth [as Carroll (1992) shows is true if consumers are behaving optimally], we can expect the approximation to be relatively good. It does not seem implausible that people could learn about two such parameters—especially since they are learning about parameters that can be given highly intuitive interpretations: \bar{X} is how much target wealth to try to have on hand, and γ indicates how quickly you try to return to that level of wealth when you are away from it.

3.2. How Good Is the Best Approximation?

The better an approximation is in utility terms, the more plausible it is that consumers would settle for the approximation rather than attempting a more exact solution. One way to measure approximation quality is to ask how much consumers who *do* know how to solve the full optimization problem would be willing to sacrifice to avoid being forced to switch permanently to the best possible approximate rule.

Answering this question requires us to define the value function associated with a rule θ . The definition is straightforward:

$$V^\theta(X_t) = E_t \left[\sum_{s=t}^{\infty} \beta^{s-t} u(C^\theta(X_s)) \right], \tag{4}$$

and it is relatively easy to compute the value of this function recursively (much easier than solving the full nonlinear optimization problem).

We can then define the “sacrifice value” as the maximum amount that a perfectly rational consumer with initial wealth X_t who is currently using the optimal rule $C^*(X)$ would be willing to pay to avoid being switched permanently to using the approximate rule $C^\theta(X)$; that is, the sacrifice value is the ϵ such that⁹

$$V^\theta(X) = V(X - \epsilon) \tag{5}$$

implying that

$$\epsilon^\theta(X) = X - V^{-1}(V^\theta(X)). \tag{6}$$

It is obvious from this equation that consumers at different levels of initial X would be willing to pay different amounts to avoid being switched. To make the sacrifice-value concept operational, it is therefore necessary to make some assumption about how consumers are distributed across levels of X when the threat to switch them to $C^\theta(X)$ occurs. Fortunately, there is a uniquely appropriate distribution to use: the ergodic distribution toward which any arbitrary initial cash-on-hand distribution will converge. [See Carroll (1997b) for a description of the methodology for calculating the ergodic distribution.] Thus, defining the ergodic cumulative distribution function for X as $F(X)$, we can calculate the average sacrifice value for a given choice of parameters θ as

$$\bar{\epsilon}^\theta = \int_{X_{\min}}^{X_{\max}} \epsilon^\theta(X) dF(X).$$

Figure 1 presents a contour plot showing “isosacrifice” contours for sacrifice values of $\bar{\epsilon} = \{0.05, 0.20, 0.40, 0.57\}$ and shows the point with the lowest sacrifice value, $\bar{\epsilon} = 0.003$, for $\theta = \{0.233, 1.243\}$.¹⁰ Note that the sacrifice value associated with $\{\gamma, \bar{X}\} = \{1, 1\}$ (the rightmost point on the horizontal axis) corresponds to the rule $C(X) = X$; the isosacrifice curve $\bar{\epsilon} = 0.57$ intercepts the horizontal axis at this point, indicating that the sacrifice value associated with the “spend everything” rule is about 0.57.

It may seem remarkable that the best sacrifice value is as low as 0.003. Figure 2 explains the mystery by showing $C^*(X)$ and $C^{\theta^*}(X)$ along with the locations of the 5th and 95th percentiles in the ergodic distribution of X under $C^*(X)$. As the figure illustrates, the linear approximation to the optimal consumption rule is quite close to the truly optimal rule over essentially the entire range from the 5th to the 95th percentiles. Furthermore, small deviations from the optimal consumption function will, by definition, result in second-order losses of utility.

4. BUFFER-STOCK SAVING AND INDIVIDUAL LEARNING

With these preliminaries out of the way, we can now turn to the central question, which is how to model the consumer’s learning process. Many models of learning




FIGURE 1. Isosacrifice contours.

in the economics literature have had a structure that can be crudely summarized as follows. The set Θ is a list of all possible actions that are available to the agent. In period t , the agent chooses a particular option, indicated by θ_t . He then observes an outcome v_t that is usually a noisy measure of the “true” value associated with choice θ_t . The agent notes the outcome and uses it in some manner to update his “beliefs” about the true value associated with that specific choice θ_t . After sampling a variety of θ ’s, the individual’s choices converge on the θ that the learning process has concluded yields the highest true value.¹¹

Our approach to individual learning will essentially follow this standard broad outline. However, in the intertemporal choice context, the primary difficulty is in a step that is assumed to be costless and immediate in most of the existing literature: Observing (even a noisy measure of) the value associated with a given choice of θ .¹²




FIGURE 2. The exact consumption rule (solid) and the best approximation (dashed).

4.1. Estimating the Value of Alternative Consumption Rules

Suppose that the consumer wishes to compare a set of potential consumption rules Θ individually designated θ_i , where, in principle, the θ_i could index alternative consumption rules of any kind (in practice, though, we will later take the θ_i to reflect alternative combinations of γ and \bar{X}). Suppose further that, for any initial level of cash-on-hand X_t , the consumer has some method by which she can make an exactly correct assessment of the expected discounted utility that each rule would yield, if used henceforth and forevermore; call this value $V^{\theta_i}(X_t)$ (we relax this assumption of observability of $V^{\theta_i}(X_t)$ momentarily). The consumer's goal is to find the θ_i that, if used forever afterward, yields the highest $V^{\theta_i}(X_t)$.

An immediate problem with this procedure is the evident possibility that the optimal θ_i could be different for different starting values X_t . If so, how would the consumer choose between two rules θ_j and θ_k , if, say, rule j performs better than rule k if $X_t = 2$ [i.e., $V^{\theta_j}(2) > V^{\theta_k}(2)$] but rule k outperforms rule j if $X_t = 3$ [$V^{\theta_k}(3) > V^{\theta_j}(3)$]? Note, however, that if one of the rules indexed by θ_i is the exactly optimal rule (i.e., the true solution to the dynamic program), the expected value yielded by that rule will exceed the expected value yielded by *any* other rule for *any* initial value of X_t , and so, the truly optimal rule would always be picked regardless of the starting X_t . Of course, if the rules indexed by θ_i do not include the exactly optimal rule, the kinds of reversals just outlined would be possible. Later, we implicitly examine the importance of this problem by having our consumers search for the optimal θ_i for several possible initial levels of wealth.

Holding initial X_t fixed for the time being, we are now in position to set forth our model of the consumer’s process for estimating the value associated with any particular θ_i . Imagine that for each θ_i the consumer forms an estimate of $V^{\theta_i}(X_t)$ by living through the experience of using that rule for n periods. That is, in period t the consumer spends $C_t = C^{\theta_i}(X_t)$, leaving $X_t - C_t$ in savings for the next period and generating period- t utility $u_t = u(C_t)$; in period $t + 1$, the consumer draws a random income shock Y_{t+1} from the distribution outlined earlier, constructs $X_{t+1} = (X_t - C_t) + Y_{t+1}$, and consumes $C^{\theta_i}(X_{t+1})$, generating period $t + 1$ utility u_{t+1} . This process is repeated until period $t + n$ is reached. As she goes, the consumer keeps track of a variable that we call “partial value”:

$$W_s = W_{s-1} + \beta^{s-t} U_s, \tag{7}$$

from a starting value of $W_{t-1} = 0$, which cumulates to

$$\begin{aligned} W_s &= U_t + \beta U_{t+1} + \beta^2 U_{t+2} + \dots + \beta^{s-t} U_s \\ &= \sum_{q=t}^s \beta^{q-t} U_q, \end{aligned} \tag{8}$$

and, when she reaches period $t + n$, she will have an estimate of the value generated by this program $\tilde{V}^{\theta_i}(X_t) = W_{t+n}$. Of course, if $n < \infty$, the value constructed in this manner will be missing a term that reflects $E_t[\beta^{n+1} V^{\theta_i}(X_{t+n+1})]$, but for n sufficiently large the omitted term should be relatively small. One purpose of our simulations is to determine the meaning of “sufficiently large” and “relatively small” in this context.

The most naive model of the individual search process would be simply to have consumers execute the foregoing procedure for a variety of potential θ_i ’s and pick the one with the highest experienced value $\tilde{V}^{\theta_i}(X_t)$. However, this procedure would produce a very noisy estimate of the true value of each possible rule because discounting by β means that the actual value experienced will be heavily influenced by the particular sequence of stochastic income draws that the consumer receives early in her experience with each rule. Even if we let n approach infinity, the consumers’ estimates of the value associated with each rule do not converge to the true values because utility from the additional later periods is discounted at an ever-higher rate and cannot overcome the initial impression made by early experience.

The only way the consumer can form a consistent estimator of the true value associated with each rule starting at X_t is to live through the experience of using each rule starting from the same X_t multiple times. That is, if the estimated value obtained the first time the consumer runs through the foregoing procedure is $\tilde{V}_1^{\theta_i}(X_t)$, the consumer will need to begin again with the same initial X_t and form a second $\tilde{V}_2^{\theta_i}(X_t)$ and so forth. We assume that the consumer runs through this

experience m times and estimates the true value of policy θ_i , starting from X_t as the average of the m experiences,

$$\hat{V}^{\theta_i}(X_t) = (1/m) \sum_{j=1}^m \tilde{V}_j^{\theta_i}(X_t).$$

It is easy to show that, as m and n jointly go to infinity, the foregoing procedure will yield an arbitrarily accurate estimate of the true value function $V^{\theta_i}(X_t)$ for any given X_t .¹³ The question that can be answered only by simulations is how large m and n need to be for the consumer to be able to have a reasonably high degree of confidence in the accuracy of her estimate $\hat{V}^{\theta_i}(X_t)$. The answer to that question, of course, depends on the metric used to evaluate \hat{V} 's accuracy. In this context, the logical metric is whether the \hat{V} 's generated by a given (m, n) combination will reliably lead the consumer to choose a good consumption rule from among the candidate rules indexed by θ_i . Before that question can be answered, however, we need to specify the process by which the set of rules to be considered is constructed.

4.2. Choosing a Set of Rules to Evaluate

Our assumption is that the θ_i simply enumerate the nodes in a grid determined jointly by the set of potentially "reasonable" values of γ and \bar{X} . For the marginal propensity to consume, the natural space of possible values is $\gamma \in [0, 1]$. Since X includes current income and the expected value of income is 1, a lower bound for \bar{X} is 1. The range of "reasonable" maximum values for \bar{X} is less obvious. Our admittedly arbitrary decision was to choose $\bar{X} \in [1, 3]$. The final assumption we need to make is about the fineness of the grid. We choose the interval between grid points for γ to be 0.05, and the interval for grid points of \bar{X} to be 0.1, for a total of $20 \times 20 = 400$ combinations of rules.¹⁴ The best of these rules is $(\gamma, \bar{X}) = (.25, 1.2)$ for which the sacrifice value is 0.007. (We are aware that a grid search is highly inefficient; we discuss robustness of our results to alternative, more efficient search procedures below.)

4.3. Results

We are now in position to specify how we will evaluate the effectiveness of various choices of m and n . We construct a population of 1,000 consumers, each of whom enters the first period of simulation with the same initial level of savings S_{t-1} (for technical reasons, this is slightly easier than starting out all consumers with the same initial values of X_t , as expositied above). For each combination of m and n , we simulate the experience of each of the 1,000 consumers executing the algorithm described above and calculating his own estimated value of \hat{V}^{θ_i} for each of the 400 possible θ_i , and, at the end of the simulations, each consumer picks the rule with the maximum estimated value $\hat{V}^{\theta_i}(X_t)$ among the rules he has tried.

Table 1 presents the results. The table is divided into three panels corresponding to different assumptions about the initial resources with which the consumers begin

TABLE 1. Search success rate and number of periods

Statistic	n^a	$m^b = 1$	$m = 10$	$m = 50$	$m = 200$
$S_{t-1} = 0$					
Mean sacrifice	10	0.289	0.197	0.160	0.135
Success ^c rate		0.07	0.17	0.18	0.13
Total periods		4,000	40,000	200,000	800,000
Mean sacrifice	20	0.218	0.110	0.073	0.060
Success rate		0.17	0.34	0.47	0.53
Total periods		8,000	80,000	400,000	1,600,000
Mean sacrifice	50	0.172	0.074	0.045	0.028
Success rate		0.24	0.46	0.68	0.86
Total periods		20,000	200,000	1.00E+06	4.00E+06
$S_{t-1} = 1$					
Mean sacrifice	10	0.269	0.122	0.100	0.102
Success rate		0.09	0.23	0.29	0.24
Total periods		4,000	40,000	200,000	800,000
Mean sacrifice	20	0.226	0.079	0.053	0.047
Success rate		0.18	0.45	0.62	0.68
Total periods		8,000	80,000	400,000	1,600,000
Mean sacrifice	50	0.187	0.058	0.036	0.024
Success rate		0.26	0.58	0.76	0.91
Total periods		20,000	200,000	1.00E+06	4.00E+06
$S_{t-1} = 2$					
Mean sacrifice	10	0.204	0.092	0.100	0.108
Success rate		0.20	0.38	0.28	0.18
Total periods		4,000	40,000	200,000	800,000
Mean sacrifice	20	0.179	0.058	0.050	0.054
Success rate		0.27	0.58	0.64	0.58
Total periods		8,000	80,000	400,000	1,600,000
Mean sacrifice	50	0.169	0.053	0.037	0.030
Success rate		0.32	0.62	0.75	0.85
Total periods		20,000	200,000	1.00E+06	4.00E+06

^a Number of periods the consumer uses a rule for each trial.

^b Number of trials.

^c Success is defined as finding a rule with sacrifice value <0.05.

the simulations, $S_0 = [0, 1, 2]$. For each (m, n) combination, three statistics are reported: the average sacrifice value of the rules picked by our 1,000 consumers, the fraction of the consumers who picked a “good” θ_i , defined as a rule with a sacrifice value of less than 5%,¹⁵ and the total number of model simulation periods each consumer has lived through in the course of searching for the rule (which will be $400 mn$).

The overwhelming conclusion from this table is that, although it is possible for this “learning by experience” method to reliably identify good consumption rules, the *amount* of experience required is staggering. The only (m, n) combination that can identify a good rule at least 80% of the time is $(m = 200, n = 50)$, which implies a search time of 4 million $(= 200 \times 50 \times 400)$ periods! Even if the criterion is merely that the (m, n) pair should produce rules with an *average* sacrifice value of 0.05 or less, the minimum number of simulation periods required is roughly a million. Interpreting the model period as a year (the appropriate interpretation for the calibration $\beta = 0.95$), it takes a million years of experience to reliably identify a reasonably good consumption rule by personal experience!¹⁶ Even reinterpreting the model period as a 2-week pay period rather than a year (a reinterpretation that is problematic because the true biweekly income process is very different from our assumed income process) leaves the required time to find a good rule absurdly long. Conclusions are roughly the same regardless of the starting values for S .

Of course, it is possible that we have not endowed our agents with enough intelligence. For instance, rather than blindly searching every point on the (γ, \bar{X}) grid, intelligent consumers could do an ordered search in which they choose a very coarse initial grid of, say, two possible choices for each of (γ, \bar{X}) , pick the best of the four choices, then center a new search grid around this optimum, and so on, or they could use a Newton algorithm or some other hill-climbing procedure. However, even if the search could be reduced so that only, say, four different rules needed to be evaluated, it would still be necessary to use values of (m, n) large enough to distinguish good rules from bad. Given that the minimum (m, n) combination that appears capable of producing the necessary accuracy is $(50, 50)$, even such a highly efficient hill-climbing routine could not reduce the number of periods required to less than $10,000 = 50 \times 50 \times 4$.

Hence, rather than alleviating the mystery of how ordinary consumers seem to have managed to learn nearly optimal consumption behavior, our exploration of the possibility of learning by experience has only deepened the mystery. On reflection, this result is not as surprising as it at first may appear. One fact that is known by any economist who has attempted numerical solution of consumption models is that finding optimal behavior in these models is an extraordinarily computationally intensive task. If there were some learning-by-experience method that could identify nearly optimal rules with vastly less computational effort, some clever economist would probably have identified that method long ago and it would now be the standard method used to solve such problems. The finding here that learning by experience requires a large amount of experience can therefore be recast as a finding that a search algorithm based on learning by experience does not drastically reduce the computational input required to find a nearly optimal rule.

5. BUFFER-STOCK SAVING AND SOCIAL LEARNING

If it takes an individual agent a million periods of experience to reliably find a good consumption rule, a population of a million consumers scattered across

the (γ, \bar{X}) landscape should *collectively* obtain essentially the same amount of information in a single period. If there were a mechanism by which all of that information could be efficiently combined, the number of model periods required for finding the optimal rule could surely be radically reduced (though computing demands remain formidable because, for each model period, calculations must now be made for many agents rather than one).

A potential mechanism to accomplish this purpose is social learning, in which individuals encounter each other and communicate the results of their own experience to others.¹⁷ Even if the social learning process is less than perfectly efficient, it still seems plausible that it might lead a population of consumers to converge on the optimum relatively quickly.

However, the existing literature on social learning has found that even fully rational social learning processes do not always result in the population as a whole reaching an optimal outcome. If each agent's actions do not fully reveal the individual's private information, the population can end up making choices little better than the choices that would be made by individual agents acting in isolation. Thus, the most interesting question to be addressed in a future literature on social learning about intertemporal choice is under what circumstances the population does and does not settle on a reasonably good set of rules.

6. CONCLUSION

Until recently, economists spent little time or effort trying to understand the cognitive processes that led to observed economic behavior, relying instead, either implicitly or explicitly, on the assumption articulated by Friedman that people could easily learn optimal behavior through trial and error. While several economic literatures have recently begun to explore the implications of learning and evolutionary dynamics, there has been little work on the role of learning in the realm of intertemporal choice.

In part, the lack of research in this area probably has been attributable to the great complexity of finding and executing optimal plans for intertemporal problems. Given this complexity, it may have seemed hopeless for consumers to learn exactly optimal behavior by experience. The first contribution of this paper is to provide an example in which the true solution can be very closely approximated (in utility terms) by a simple linear model of behavior in which both slope and intercept have intuitive (and plausibly learnable) interpretations.

The second contribution of the paper is to show that even when the goal is to learn only this simple approximation, pure trial-and-error learning requires an enormous amount of experience to allow consumers to distinguish good rules from bad ones—far more experience than any one consumer would have over the course of a single lifetime.

These results suggest that the learning model proposed here is not an adequate description of the process by which consumers learn about consumption behavior. It remains possible, of course, that consumers employ an individual learning

mechanism that is much more efficient than the one postulated here, and the search for an improved individual learning algorithm is a possible direction for future research. Such a learning mechanism, if found, should also constitute an important advance in the technology for solving dynamic optimization problems.

More intriguing, however, is the possibility that consumers come by their behavior by a process of social learning, in which rules of thumb that are successful in utility terms are passed along from one consumer to another, or through other mechanisms such as the advice of personal finance experts or advice in personal finance books. In fact, personal finance books often give advice that sounds very much like buffer-stock saving behavior with respect to liquid assets [see Carroll (1997a) for a typical reference in a personal finance book]. Elucidating the circumstances under which a process of social learning can be expected to lead the population to reasonably optimal behavior will be an interesting task for future work.

NOTES

1. The principal example of which we are aware is Lettau and Uhlig (1999); Marcet and Sargent (1991) examine an investment problem but assume that decision makers know dynamic optimization theory and are only learning about the distribution of shocks.

2. See Deaton (1991), Carroll (1997a), Cagetti (1999), Gourinchas and Parker (1999), and the papers cited therein for evidence.

3. Perhaps the most common test of this kind has been in the context of determining the effects of Social Security and of other defined benefit pension schemes on personal saving. See Carroll (1994, 1997a) for other examples.

4. Carroll (1996) shows that the relevant condition is $R\beta E_t(G\tilde{N}_{t+1})^{-\rho} < 1$, where R is the interest rate, β is the time preference factor, G is the growth rate of income, ρ is the coefficient of relative risk aversion, and N is the mean-one multiplicative shock to permanent income. Parameter values used by Carroll (1997a) were a time preference rate of 4% annually, household income growth of 3%, coefficient of relative risk aversion of 3, and a real after-tax interest rate of 0%; results were robust to plausible variation in these parameters.

5. Hubbard et al. (1994, 1995) had to use a supercomputer to solve the optimal life-cycle problem when it was enhanced to incorporate a modest degree of realism about health and mortality risk and the structure of social insurance programs.

6. For example, the limiting MPC as wealth goes to infinity or zero can be calculated [see Carroll (1996)], and Carroll and Kimball (1996) prove that the consumption rule is strictly concave.

7. See note 4 for the condition (where, here, $N = 1$ because we have assumed that there are no permanent shocks).

8. We consider a model with explicit liquidity constraints here because it is somewhat simpler than the model without constraints. Carroll (2001) shows that, in many important respects, the optimization models with and without constraints are essentially the same because the precautionary saving motive serves as a kind of self-imposed liquidity constraint.

9. This definition is inspired by the definition of the equivalent risk premium from consumption theory.

10. The θ that minimizes the sacrifice value is not exactly the same as the $\theta = \{\gamma^*, \bar{X}^*\}$, which constitutes the first-order Taylor approximation to the optimal rule, though the two θ 's are close.

11. This summary does not encompass what might be termed "general equilibrium" learning problems in which the average choices of the collection of agents actually affect the payoff that each individual choice provides.

12. Several previous authors in the macroeconomics literature have assumed that consumers understand dynamic stochastic optimization theory and that their “learning problem” is to discover the properties of the stochastic processes that impinge on their optimization problem [see, e.g., Sargent (1993, pp. 93–107) and Marcet and Sargent (1991)]. The only paper we are aware of that examines agents’ ability to learn how to solve a true dynamic stochastic optimization problem is that of Lettau and Uhlig (1999), who examine the (in)ability of artificial intelligence constructs called “classifier systems” developed by Holland (1986) to learn the dynamic programming solution to an optimal consumption problem. Unfortunately, to use Holland’s classifier systems, Lettau and Uhlig must drastically reduce the complexity of the optimal consumption problem. Their central example is one in which consumers have a choice of only two possible levels of consumption, and there are only three possible levels of wealth. They find that a rule of thumb of always spending all available resources is not driven out of the classifier system in the long run, essentially because the mechanism for updating the strengths of the different classifier rules does not correspond to the prescriptions of dynamic stochastic optimization. Although this is an important and interesting paper, it appears to have little relevance to the approach we pursue here.

13. We have verified that the estimates of the value obtained for very large values of m and n are extremely close to the estimates obtained through our completely independent theoretical exercise of constructing the value function directly.

14. We exclude the value $\gamma = 0$ from the set under consideration because all rules with $\gamma = 0$ are identical regardless of the value of \bar{X} . This is why there are 20 rather than 21 possible values of γ . To obtain 20 rather than 21 values of \bar{X} , we exclude $\bar{X} = 3.0$ from the list.

15. Out of a total of 400 rules, there were 20 for which the sacrifice value was less than 5%.

16. Note that this assumes that consumers do not need to explore alternative starting values for S_{t-1} . If we were to assume that they search over three values of S_{t-1} as presented in the table, search times would triple.

17. We know of only two papers that present any empirical evidence on social learning and saving behavior. Lusardi (1999) finds that consumers are more likely to have thought about and prepared for retirement if they have an older sibling who has already retired. Bernheim (1996) finds that workplace retirement education courses changed worker behavior with respect to 401(k) retirement accounts in directions that most economists would identify as more rational.

REFERENCES

- Arthur, W.B., J.H. Holland, B. LeBaron, R. Palmer & P. Tayler (1997) Asset pricing under endogenous expectations in an artificial stock market. In W.B. Arthur, S.N. Durlauf, & D. Lane (eds.), *The Economy as an Evolving Complex System II*. Reading, MA: Addison Wesley Longman.
- Bernheim, D. (1996) Personal saving, information, and economic literacy: New directions for public policy. In *Tax Policy for Economic Growth in the 1990s*, pp. 53–78. Washington, DC: American Council for Capital Formation.
- Bikhchandani, S., D. Hirshleifer & I. Welch (1998) Learning from the behavior of others: Conformity, fads, and informational cascades. *Journal of Economic Perspectives* 12(3), 151–170.
- Browning, M.J. & A. Lusardi (1996) Household saving: Micro theories and micro facts. *Journal of Economic Literature* 34, 1797–1855.
- Cagetti, M. (1999) Wealth Accumulation Over the Life Cycle and Precautionary Savings. Manuscript, University of Chicago.
- Carroll, C.D. (1992) The buffer-stock theory of saving: Some macroeconomic evidence. *Brookings Papers on Economic Activity*, 1992(2), 61–156.
- Carroll, C.D. (1994) How does future income affect current consumption? *Quarterly Journal of Economics* 109, 111–148 (<http://www.econ.jhu.edu/people/ccarroll/howdoesfuture.pdf>).
- Carroll, C.D. (1996) Buffer Stock Saving: Some Theory. Manuscript, Johns Hopkins University (<http://www.econ.jhu.edu/people/ccarroll/BufferTheory.pdf>).

- Carroll, C.D. (1997a) Buffer-stock saving and the life cycle/permanent income hypothesis. *Quarterly Journal of Economics* 112, 1–56 (<http://www.econ.jhu.edu/people/ccarroll/BSLCPH.pdf>).
- Carroll, C.D. (1997b) Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second Order Approximation). NBER working paper 6298 (<http://www.econ.jhu.edu/people/ccarroll/death.pdf>).
- Carroll, C.D. (2001) (forthcoming) A theory of the consumption function, with and without liquidity constraints. *Journal of Economic Perspectives*.
- Carroll, C.D. & M.S. Kimball (1996) On the concavity of the consumption function. *Econometrica* 64, 981–992 (<http://www.econ.jhu.edu/people/ccarroll/concavity.pdf>).
- Deaton, A.S. (1991) Saving and liquidity constraints. *Econometrica* 59, 1221–1248.
- Deaton, A.S. (1992) *Understanding Consumption*. New York: Oxford University Press.
- Ellison, G. & D. Fudenberg (1995) Word-of-mouth communication and social learning. *Quarterly Journal of Economics* 110, 93–125.
- Evans, G.W. & S. Honkapohja (1999) Learning Dynamics. Manuscript, University of Oregon.
- Friedman, M.A. (1953) *Essays in Positive Economics*. Chicago: University of Chicago Press.
- Gale, D. (1996) What have we learned from social learning? *European Economic Review* 40, 617–628.
- Gourinchas, P.-O. & J. Parker (1999) Consumption over the life cycle. Manuscript, Princeton University.
- Hall, R.E. (1978) Stochastic implications of the life-cycle/permanent income hypothesis: Theory and evidence. *Journal of Political Economy* 96, 971–987.
- Hall, R.E. & F. Mishkin (1982) The Sensitivity of consumption to transitory income: Evidence from PSID households. *Econometrica* 50, 461–481.
- Holland, J.H. (1986) *Adaptation in Natural and Artificial Systems*. Cambridge, MA: MIT Press.
- Hubbard, R.G., J.S. Skinner & S.P. Zeldes (1994) The importance of precautionary motives for explaining individual and aggregate saving. *Carnegie-Rochester Conference Series on Public Policy* 40, 59–126.
- Hubbard, R.G., J.S. Skinner & S.P. Zeldes (1995) Precautionary saving and social insurance. *Journal of Political Economy* 103, 330–399.
- Lettau, M. (1997) Explaining the facts with adaptive agents: The case of mutual fund flows. *Journal of Economic Dynamics and Control* 21, 1117–1147.
- Lettau, M. & H. Uhlig (1999) Rules of thumb and dynamic programming. *American Economic Review* 89, 148–174.
- Lusardi, A. (1999) Explaining Why So Many Households Do Not Save. Manuscript, Dartmouth College.
- Marcet, A. & T.J. Sargent (1991) Convergence of least squares learning mechanisms in self referential linear stochastic models. *Journal of Economic Theory* 48, 337–368.
- Marimon, R. (1997) Learning from learning in economics. In D. Kreps & K. Wallis (eds.), *Advances in Economics and Econometrics: Theory and Applications*, Vol. I. Cambridge, UK: Cambridge University Press.
- McCarthy, J. (1995) Imperfect insurance and differing propensities to consume across households. *Journal of Monetary Economics* 36, 301–327.
- Parker, J. (1999) The reaction of household consumption to predictable changes in social security taxes. *American Economic Review* 89, 959–973.
- Samuelson, L. (1997) *Evolutionary Games and Equilibrium Selection*. Cambridge, MA: MIT Press.
- Sargent, T.J. (1993) *Bounded Rationality in Macroeconomics*. Oxford, Oxford University Press.
- Soules, N.S. (1999) The response of household consumption to income tax refunds. *American Economic Review* 89, 947–958.
- Young, H.P. (1997) *Learning and Evolution in Games*. Princeton, NJ: Princeton University Press.
- Zeldes, S.P. (1989) Optimal consumption with stochastic income: Deviations from certainty equivalence. *Quarterly Journal of Economics* 104, 275–298.