Very Preliminary
Comments Welcome

Learning About Intertemporal Choice

Todd Allen[†]
twa@lacasa.com

Christopher D. Carroll[‡]
ccarroll@jhu.edu

May 21, 1999

[‡] NBER and The Johns Hopkins University. Correspondence to Christopher Carroll, Department of Economics, Johns Hopkins University, Baltimore, MD 21218-2685 or ccarroll@jhu.edu.

[‡] Center for Adaptive Systems Applications 1302 Osage, Suite A, Santa Fe, NM 87504

1 Introduction

Economists long ago adopted the convention that economic behavior should be modelled 'as if' the economic agents were solving a formal mathematical optimization problem. While this methodology has never been entirely uncontroversial, and in the game theory literature has recently been challenged by theories based on learning and evolutionary dynamics (for a summary see, e.g., Young (1997) or Samuelson (1997)), mathematical optimization remains virtually the only modelling paradigm in the literature on household consumption and portfolio choice behavior,[1] as well as in many other branches of economics outside of game theory.[2]

Perhaps the central intellectual underpinning of the rational optimization approach in the consumption literature (and often elsewhere) is Milton Friedman's (1957) famous analogy to the game of pool. Friedman noted that a scientist wanting to model the behavior of a pool player would find it essential to assume the player could solve the sophisticated problems in Newtonian physics reflected in the movements of the balls. Friedman argued that, while the player was not solving formal mathematical problems in his head, experience could give him adequate intuition about the relevant physics. Similarly, Friedman argued, consumers could be expected eventually to learn the optimal solution to the repeated problem of deciding how much of their income to consume.

When Friedman proposed his 'learning hypothesis' as the key underpinning for the 'as if' optimization paradigm, the state of knowledge about human cognition and learning was far too primitive to allow for any rigorous examination of the hypothesis. Today, however, advances in mathematical techniques and computer speed, along with important developments in the game theory, evolutionary biology, cognitive science, and artificial intelligence literatures promise to provide the tools and ideas with which we can begin to explore the 'learning hypothesis' rigorously. The central purpose of the research program described in this proposal is to begin that exploration in the context of consumption, portfolio choice, and other intertemporal choice behavior.

2 Background and Literature Summary

Unfortunately for economists, the optimal saving/consumption problem does not have an analytical solution under plausible specifications of utility and uncertainty; as a substitute, until

[1] Thaler (1994), Shiller (1997) and others have attempted to bring insights from behavioral psychology to bear on saving and portfolio decisions, but as yet this strategy has not caught on widely in the literature, in part because there is not a well-defined and articulated framework underlying the often compelling individual points these authors make.

[2] Learning models have also made a modest appearance in macroeconomic theory, particularly with application to the problem of deciding among a multiplicity of rational expectations equilibria. Though these issues are closely related to the problem of multiple equilibria in game theoretic contexts, an important difference is that there is usually no strategic element in the behavior of individuals in macroeconomic models. See Sargent (1993) for an excellent summary of this literature.

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.

From the 1950s to the present, this Certainty Equivalent (CEQ) model has been tested exhaustively. A recent summary of the literature suggests that the model fails in at least three ways (Deaton (1992)). First, a large literature dating from the 1950s and 1960s and extending through Hall and Mishkin (1982) and Souleles (1995), consistently estimated a marginal propensity to consume out of transitory income greater than 0.2. Since the CEQ model generally implies MPC's of less than 0.05, these results were interpreted as suggesting the presence of some consumers who are either rational but liquidity constrained or simply always consume their income. Second, another large literature 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 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 1960s to the most recent (1995) triennial Survey of Consumer Finances, have found that the median household at every age before about 50 typically holds total non-housing 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 solutions of the optimal consumption problem under realistic assumptions about uncertainty and preferences,[4] all of these empirical findings 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 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 percent or higher, the propensity to consume out of human wealth can be close to zero, and standard Euler equation tests of

[3] Perhaps the commonest 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 in Carroll (1997a) were a time preference rate of 4 percent annually, household income growth of 3 percent, coefficient of relative risk aversion of 3, and a real after-tax interest rate of 0 percent; results were robust to plausible variation in these parameters.

consumption behavior 'fail' in ways that can replicate the whole range of empirical failures of the Euler equation. (See Carroll (1992, 1997a, 1997b) for details). Uncertainty and the consequent precautionary saving motive thus turn out to profoundly modify optimal behavior from that which is predicted by the model that was taken by economists to embody "rationality" from the 1950s through the late 1980s.

In a way, the recent findings can be interpreted as a potential vindication of Friedman's argument that people can grasp the solution to a difficult mathematical problem even without mathematical training. The embarrassment is that economists for so long failed to see what consumers apparently implicitly know:[5] that buffer-stock saving behavior works reasonably well. But these findings also raise rather urgently the question of how ordinary consumers appear to be able to solve, even approximately, problems that even now, and even in versions much simpler than the actual problems people face, continue to strain the capabilities powerful modern computers.[6] One possible answer is that people may have a powerful inbuilt intuition about the solution to dynamic optimization problems. But 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, since 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, but it is an explanation whose credibility would be considerably greater if it had ever been seriously examined and tested – as this project proposes to do.

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 of liquid assets which they use to smooth consumption in the face of an uncertain income stream. If their buffer stock of precautionary assets falls below the target, they will consume less than their expected income and liquid assets will rise, while 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. While certain analytical characteristics of the rule can be proven,[7] it is hard to see how a consumer without a supercomputer and a Ph.D. could be expected to determine the *exact*

[5] Perhaps from personal finance books. See Carroll (1997a) for a typical reference from a personal finance book.

[6] Hubbard, Skinner, and Zeldes (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.

[7] 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.

shape of the nonlinear and nonanalytical decision rule.

Fortunately for consumers, it turns out not to matter much whether they get the fine details of the rule right: Simulation experiments show that simple and intuitive approximations to the optimal rule can generate utility streams that are only modestly less 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 $\rho = 3$.[8] Income is stochastic with a 3-point distribution $(.7, 1, 1.3)$ with probabilities $(.2, .6, .2)$, a process chosen to (very roughly) match empirical evidence on the amount of transitory variation in annual household income observed in the PSID (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 $\beta = .95$. The traditional approach to modelling consumer behavior is to suppose that the consumer solves the problem:

$$\begin{aligned} \max_{\{C_s\}} \quad & E_t \sum_{s=t}^{\infty} \beta^{s-t} u(C_s) \\ \text{s.t.} \quad & X_{s+1} = X_s - C_s + \tilde{Y}_{s+1} \\ & X_s \geq 0 \;\; \forall \;\; s \end{aligned} \tag{1}$$

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

$$V(X_t) = \max_{\{C_s\}} \quad u(C_s) + \beta V(X_{t+1}) \tag{2}$$

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 X^*. Formally, Carroll (1996) shows that if the parameters of the problem satisfy a certain 'impatience' condition[9] then an X^* will exist such that if $X_t > X^*$ then $E_t X_{t+1} < X_t$ and if $X_t < X^*$ then $E_t X_{t+1} > X_t$. Furthermore, because expected income is 1 (and X_{t+1} includes Y_{t+1}), if consumption is less than 1 we know that $E_t X_{t+1} > X_t$ and if consumption is greater than 1 we know that $E_t X_{t+1} < X_t$. Assuming $X^* \geq 1$, the optimal consumption rule can be rewritten, without loss of generality, as:

$$C^*(X) \;\; = \;\; 1 + f(X - X^*). \tag{3}$$

Using the fact that $E_t \tilde{Y}_{t+1} = 1$ we know that $E_t X_{t+1} = X_t - C_t + 1$. But at the point where $X = X^*$ we have $E_t X_{t+1} = X_t$ which implies that $X_t - C_t + 1 = X_t$ which implies that

[8] In order to prevent the Inada condition from unduly influencing outcomes, we also assume that consumption never falls below ten percent of permanent income. For simplicity we leave this constraint out of the treatment in the text.

[9] See footnote 3 for the condition.

$C_t = 1$. Hence we know that $f(0) = 0$. Calling $\gamma = f'(0)$, a first-order Taylor expansion of equation (3) around the point $X = X^*$ is therefore

$$\tilde{C}(X) \approx 1 + \gamma(X - X^*). \tag{4}$$

To capture the liquidity constraint, suppose actual consumption is given by:

$$\hat{C}(X) = \begin{cases} \tilde{C}(X) & \text{if } \tilde{C}(X) \leq X, \\ X & \text{if } \tilde{C}(X) \geq X. \end{cases}$$

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, X^* and γ. This kind of simple rule seems at least roughly consistent with the wisdom from the experimental game theory literature about the nature of the behaviors people appear capable of learning (Roth and Erev (1995)) – especially since they are learning about parameters that can be given highly intuitive interpretations: 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.

The better this approximation is in utility terms, the more plausible it is that consumers would settle for such an 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.[10] The answer turns out to be that a set of consumers behaving according to the optimal rule would on average be willing to sacrifice an amount equal to less than 1/2 percent of their average consumption to avoid being switched over to the optimal approximate rule (henceforth for brevity we will refer to the amount that optimal consumers would pay to avoid being switched to an approximate consumption rule as the 'sacrifice value' associated with that approximate rule).[11] Hence, the two-parameter piecewise linear approximation can allow shrewd consumers to do nearly as well as the full model, and so it seems plausible to model the consumer's problem as being that of 'learning' the optimal values of γ and X^*.

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. Several previous authors in the macroeconomics

[10]Because consumers at different levels of X would be willing to pay different amounts to avoid being switched, the answer will depend on an assumption about how consumers are distributed across different levels of cash-on-hand. 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).

[11]This calculation was performed essentially by numerically searching for the (γ, X^*) combination that minimizes the sacrifice value. This function is single-peaked and well-behaved. The approximate rule with the lowest sacrifice value is $(\gamma, X^*) = (.236, 1.24)$ with a sacrifice value of 0.004.

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 game theory literature has explored a wider array of assumptions about learning processes, but many of the ideas explored in that literature are adapted to learning about strategic interactions with other agents, which is not relevant in the current context. Our initial approach is minimalist: we assume that our consumers observe only the stream of income and utility realizations that result when they pursue a given policy; they know neither the stochastic structure of the shocks nor how to solve dynamic stochastic optimization problems. A crude statement of the initial question to be answered is whether such consumers can find a good approximation to the optimal rule simply by 'trying out' a variety of potential rules and then choosing those that perform best. This methodology is related to the literature on reinforcement learning in which agents experiment with various behaviors and then increasingly recur to those behaviors which have yielded high payoffs in the past (see Arthur (1993) for an argument that this is a good description of actual human behavior, and Roth and Erev (1995) for supportive evidence). However, our agents are somewhat more purposive than reinforcement learners often are: we assume that agents systematically explore the entire space of possible behaviors and then adopt the best strategy among those examined.

4.1 Estimating the Value of Alternative 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 (though in practice we will later take the θ_i to reflect alternative combinations of γ and 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 each rule would yield, if used; call this value $V^{\theta_i}(X_t)$ (we will relax this assumption of perfect observability of $V^{\theta_i}(X_t)$ momentarily). Then for any possible X_t the consumer's problem would simply be to pick the θ_i for which $V^{\theta_i}(X_t)$ is greatest.

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 $X_t = 2$ but rule k outperforms rule j if $X_t = 3$? Note, however, that if one of the rules indexed by θ_i is the exactly optimal rule, 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. Below we will implicitly examine the importance of this problem by having our consumers search for the optimal θ_i for several possible initial X_t.

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, designating the consumption rule associated with θ_i as $\hat{C}^{\theta_i}(X)$, 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 $\tilde{Y}_{t+1}$ from the distribution outlined above, constructs $X_{t+1} = (X_t - C_t) + \tilde{Y}_{t+1}$, and consumes $\hat{C}^{\theta_i}(X_{t+1})$, generating period $t+1$ utility u_{t+1}. This process is repeated until period $t+n$ is reached. As she goes, the consumer updates her estimate of the value generated by this program $\tilde{V}^{\theta_i}(X_t)$ by cumulating and discounting the period-utility functions appropriately. 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' 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 the actual value experienced will be heavily influenced by the particular sequence of stochastic income draws 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, $\hat{V}^{\theta_i}(X_t) = (1/m) \sum_{j=1}^{n} \tilde{V}_j^{\theta_i}(X_t)$.

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.[12] 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

[12]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.

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 possible rules θ_i is constructed.

4.2 Choosing a Set of Rules to Evaluate

Our initial 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 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 X^* is 1. The range of 'reasonable' maximum values for X^* is less obvious. Our preliminary and admittedly arbitrary decision was to choose $X^* \in [1, 3]$. 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 X^* to be 0.1, for a total of 20x20=400 combinations of rules.[13] The best of these rules is $(\gamma, X^*) = (.25, 1.2)$ for which the sacrifice value is 0.008.

4.3 Very Preliminary Results

We are now (finally) in position to specify how we will evaluate the effectiveness of various choices of m and n. We construct a population of 100 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 exposited above). For each combination of m and n we simulate the experience of each of the 100 consumers executing the alogrithm described above and calculating their own estimated value of $\hat{V}^{\theta_i}$ 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 the simulations, $S_0 = [0, 1, 2]$. For each (m, n) combination, three statistics are tabulated: the average sacrifice value of the rules picked by our 100 consumers, the fraction of the consumers who picked a 'good' θ_i, defined as a rule with a sacrifice value of less than 5 percent,[14] and the total number of model simulation periods each consumer has lived through in the course of searching for the rule (which will be 400mn).

The overwhelming conclusion from this table is that, while 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) combinations that can identify a good rule at least 80 percent of the time is $(m = 200, n = 50)$ which implies a serch time of 4 million (=200*50*400)! Even if the criterion is merely that the (m, n) pair should produce rules with

[13] We exclude the value $\gamma = 0$ from the set under consideration because all rules with $\gamma = 0$ are identical regardless of the value of X^*. This is why there are 20 rather than 21 possible values of γ. In order to obtain 20 rather than 21 values of X^* we exclude $X^* = 3.0$ from the list.

[14] Out of a total of 400 rules, there were 17 for which the sacrifice value was less than 5 percent.

Individual Search Results

		$m = 1$	$m = 10$	$m = 50$	$m = 200$
		$S_{t-1} = 0$			
$n = 10$	Mean Sacrifice:	0.275	0.194	0.152	0.1388
	Success Rate:	0.07	0.12	0.14	0.16
	Total Periods:	4000	40000	200000	800000
$n = 20$	Mean Sacrifice:	0.212	0.103	0.081	0.0546
	Success Rate:	0.18	0.36	0.42	0.52
	Total Periods:	8000	80000	400000	1600000
$n = 50$	Mean Sacrifice:	0.166	0.077	0.051	0.0304
	Success Rate:	0.25	0.39	0.62	0.83
	Total Periods:	20000	200000	1.00E+06	4.00E+06
		$S_{t-1} = 1$			
$n = 10$	Mean Sacrifice:	0.297	0.115	0.099	N/A†
	Success Rate:	0.08	0.22	0.23	N/A†
	Total Periods:	4000	40000	200000	N/A†
$n = 20$	Mean Sacrifice:	0.231	0.063	0.064	N/A†
	Success Rate:	0.14	0.55	0.43	N/A†
	Total Periods:	8000	80000	400000	N/A†
$n = 50$	Mean Sacrifice:	0.198	0.07	0.045	N/A†
	Success Rate:	0.29	0.45	0.66	N/A†
	Total Periods:	20000	200000	1.00E+06	N/A†
		$S_{t-1} = 2$			
$n = 10$	Mean Sacrifice:	0.177	0.081	0.09	N/A†
	Success Rate:	0.23	0.42	0.24	N/A†
	Total Periods:	4000	40000	200000	N/A†
$n = 20$	Mean Sacrifice:	0.161	0.059	0.052	N/A†
	Success Rate:	0.27	0.53	0.58	N/A†
	Total Periods:	8000	80000	400000	N/A†
$n = 50$	Mean Sacrifice:	0.177	0.045	0.037	N/A†
	Success Rate:	0.28	0.67	0.73	N/A†
	Total Periods:	20000	200000	1.00E+06	N/A†

n is the number periods the consumer uses a rule for each trial.
m is the number of trials
'Success' is defined as finding a rule with sacrifice value < 0.05.
† We were unable to complete the simulations for $m = 200$
and $S_{t-1} = (1, 2)$ in time for inclusion in this table.

Table 1: Search Success Rate and Number of Periods

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 $\beta = .95$), it takes a million years of experience to reliably identify reasonably good consumption rule by personal experience![15] Even reinterpreting the model period as a two-week pay-period rather than a year (a reinterpretation that is problematic for reasons detailed below) 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, X^*) grid, intelligent consumers could do an ordered search in which they choose a very coarse initial grid (say, two possible choices for (γ, X^*), pick the best of the four choices, then center a new search grid around this optimum, and so on. Alternatively, we could assume that consumers are smart enough to immediately rule out the most extreme values of γ and X^* that we now allow them to consider. We intend to explore each of these possibilities, and others, as the project progresses. But even if the grid search could be reduced to, say, 4 binary choices, 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 grid search could not reduce the number of periods required to less than $40000 = 50 * 50 * 2^4$.

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 that mystery. Returning to Friedman's example of the pool player, what seems glaringly obvious in retrospect is that the pool player can practice hundreds of shots, and observe thousands of interactions of balls, in the course of a single day of practice. Consumers, by contrast, only experience major income shocks on an occasional basis. If consumers were able to accumulate as much experience with their problem as the pool player can accumulate with his, then they might similarly be expected to be able to learn the approximately correct answers by experience. But with large income shocks as infrequent as they are in reality, prospects of learning optimal consumption behavior by personal experience seem dim.

Personal experience, however, is not the only kind of information that consumers might be able to bring to bear: It seems reasonable to suppose that consumers may exchange information on their experiences with friends, family, colleagues, and others. This raises the question of whether, by combining the communicated experience of others with their own experience, consumers might be able to find good consumption rules much more quickly than they could on their own.

[15] 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.

5 Buffer Stock Saving and Social Learning

No man is an island - John Donne, *Devotions upon Emergent Occasions*, no. 17 (1624).

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, 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, then the search time for finding the optimal rule could surely be radically reduced.

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

Our model of social learning is related to several strands of work in the game theory literature, most notably work on imitation dynamics in Binmore and Samuelson (1997) and the work summarized by Weibull (1995). Our model also draws directly on the extensive literature on 'genetic algorithms' pioneered by evolutionary biologists and computer scientists. This literature studies the evolution of populations of artificial agents scattered across a 'fitness landscape' which stochastically maps the agents' inherent characteristics (their 'genotype') into a probability of reproductive success or failure. In our reinterpretation, the various possible consumption rules Θ constitute the analogy to a collection of genotypes, and the tendency of successful rules to be transmitted to new consumers and unsuccessful ones to be discarded corresponds to the differential 'reproductive success' of different genotypes/rules. We will discuss other relationships to the literature on genetic algorithms and broader topics in the concluding section.

5.1 Implementation

Consider a consumer j arriving in period s who began using his current rule in period $t < s$. This consumer draws a random income shock $\tilde{Y}_{j,s}$ from the distribution described above; this income combines with preexisting savings to yield $X_{j,s}$ which, from the consumer's location in (γ, X^*) space, determines consumption $C_{j,s}$ (and therefore utility $U_{j,s}$).

In order to be able to compare the performance of his rule with the performance of others' rules, the consumer needs to keep a running estimate of how well his rule has worked. If he were to encounter only consumers who had been using their rules exactly as long as he has been using his, the obvious method of tabluation would work: the consumer could simply keep track of a variable we will call 'partial value'

$$W_{j,s} = W_{j,s-1} + \beta^{s-t} U_{j,s}, \tag{5}$$

which would cumulate to

$$\begin{aligned} W_{j,s} &= U_{j,t} + \beta U_{j,t+1} + \beta^2 U_{j,t+2} + \ldots + \beta^{s-t} U_{j,s} \\ &= \sum_{q=t}^{s} \beta^{q-t} U_{j,q}. \end{aligned} \tag{6}$$

Note that, if the consumer expected to live with the current rule forever, the expectation as of time t of $W_{j,s}$ as $s \rightarrow \infty$ would be equal to the value function $V^{\theta_i}(X_t)$.

Of course, for any useful social learning to take place, consumers must evaluate and compare their rules before $s = \infty$. Suppose for discussion that all consumers use their rules for exactly n periods before comparing their rules with those of others. In this case, the expectation as of time t of 'partial value' is $E_t W_{j,t+n} = V^{\theta_j}(X_t) - \beta^n E_t V^{\theta_j}(X_{t+n})$. Truncating the consumer's experience at n periods is thus equivalent to giving the consumer a discontinuous discounting function in which the future more than n periods away is discounted at a 100 percent rate. Intuitively, this implies that, compared to the optimum rule, there is likely to be some rule that results in higher consumption over the next n periods but lower consumption thereafter whose expected partial value will be greater than the expected partial value of the optimal rule. In other words, cutting off the consumer's experience after n periods will tend to bias upward the consumer's estimate of the value of relatively 'impatient' rules. The magnitude of any such bias can be examined by simulating the model for a variety of values of n.

Now consider relaxing the constraint that all consumers who encounter each other must have lived exactly n periods with their respective rules; instead, we want to allow consumers to compare and swap rules with each other so long as both consumers have had *at least* n periods of experience with their rules. But how should two consumers who have been using their rules for different durations compare their experiences? Our assumption is that the consumers each renormalize their W's in a way that puts the two partial values on a common footing. Specifically, a consumer who has lived with his current rule for n periods would have a scaled measure (denoted by the lower-case letter)

$$w_{j,s} = \frac{W_{j,s}}{1 + \beta + \beta^2 + \ldots + \beta^n}. \tag{7}$$

In effect, the w that results from this rescaling is a weighted average of the level of utility experienced by consumer j since he began using his current rule, where the weight for each period's utility corresponds to its geometric discounting factor relative to the period in which the rule was adopted.

Now consider the encounter between consumer j and consumer k, each of whom has been using his current rule for at least n periods. From the vantage point of either consumer, there are two kinds of possible errors: adopting the other consumer's rule when it is actually inferior to one's own, or failing to adopt the other's rule when it is actually superior. If the expected value of every rule could be observed or calculated directly, consumers would make neither of these mistakes and convergence of the population to the optimal rule would be very rapid. However, given that $w_{j,s}$ and $w_{k,s}$ are imperfect measures of the value of j and k's respective rules, it is important to think through the set of circumstances in which consumers are likely to err, and the consequences of those errors.

Suppose for discussion that $w_{j,s} > w_{k,s}$, and consider k's decision whether to adopt j's rule. There are at least two obvious ways in which $w_{j,s}$ could exceed $w_{k,s}$ even if j's rule is actually inferior to k's: j might have started out with a higher level of initial resources, or might have drawn a higher level of random income shocks. Our baseline adoption rule is therefore that k adopts j's rule only if j's average weighted income[16] and initial resources are both less than or equal to k's.

Even with these restrictions, there remains at least one way in which $w_{j,s}$ could be greater than $w_{k,s}$ even if k's rule is *ex ante* superior: consumer j might have been luckier than k not in the mean *level* of income but in the *volatility* of income, given its weighted mean level. Income volatility will partially translate into consumption volatility, giving k a lower value of $w_{k,s}$ than would have obtained had k drawn j's shocks. This set of circumstances is potentially problematic because it can lead to a downward bias in the socially learned value of X^*. To see why, note that a relatively imprudent rule under which the consumer spends down his wealth leaving little or no buffer stock will yield higher w than a more prudent rule *so long as the bad shocks do not occur.* Thus the consumers who are imprudent but lucky will have a higher value of w than those who are prudent but lucky, and so the imprudent rules may spread more than they should.

In this same encounter, by assumption j will not adopt k's rule because $w_{j,s} > w_{k,s}$. Yet this could also be an error on j's part, because j's higher utility could just reflect better luck on j's part in the early period following adoption of his rule (recall that the discounting of utility makes experiences increasingly irrelevant as the distance from rule adoption date increases). We found in preliminary simulations that a small number of agents with objectively bad (imprudent) rules but who drew an exceptionally lucky string of early income shocks were essentially immovable from their bad rules. To deal with this problem, we adopt a somewhat ad-hoc solution: any consumer who reaches his 200th period using the same rule is arbitrarily assigned some other rule drawn randomly from the population of rules in use at that time. We intend to experiment with other ways of dealing with this problem, but have not had time to do so.

5.2 Very Preliminary Results

As in the individual learning simulations, we again divide the (γ, X^*) landscape into a grid with 20 distinct values of $\gamma \in [0.05, 0.10, \ldots 1.00]$ and $X^* \in [1.0, 1.1, \ldots, 2.9]$. We begin our simulation in period 1 with a population of consumers distributed evenly across this landscape, with 10 consumers initially residing at each possible rule. Table 2 reports a

[16] The weighting is identical to that in equation (7) for W.

variety of statistics that summarize the state of affairs in our population as the simulation evolves. First, we report the population average values of (γ, X^*). Second, as a more precise measure of how close the population as a whole is to optimal behavior, we calculate the sample mean of the sacrifice value (the pattern is largely the same for medians). Third, we report the proportion of the 400 original rules that remain in use by at least one member of the population.

In comparison with the individual learning case, these results confirm the intuition that social learning can greatly speed the process of finding a reasonably good solution: for $n = 20$, within about 500 simulation periods the average consumer is using a rule that has a sacrifice value of only about 4 percent. As expected, there is some tradeoff between n, the number of periods of experience required before social interchange is allowed, and bias in the socially learned solution: Requiring greater experience slows convergence but improves the accuracy of the final solution. After 10,000 periods of simulation, for example, the average sacrifice value for the $n = 20$ case is only 0.0127, compared to 0.0198 for $n = 10$. However, increasing n beyond 20 periods does not result in further improvement in the eventual solution; the sacrifice value after 10,000 periods in the $n = 50$ case is essentially identical to the sacrifice value in the $n = 20$ case after the same number of periods, and the value of X^* is only trivially higher. Interestingly, however, the $n = 10$ case produces better results than $n = 20$ or $n = 50$ for the first 100 periods of simulation, simply because the $n = 10$ case is able to eliminate the worst rules more quickly.

While the speed with which good rules are identified is greatly improved over the individual learning case, even the swiftest of the social learning configurations must run over four hundred periods (not reported in the table) before the average sacrifice value drops below our threshhold value of 5 percent. If the model period is interpreted as a year, this is still unrealistically slow.

One factor that inevitably slows convergence is our assumption that consumers simply ignore any information coming from anyone who was luckier in either initial endowment or income draws. Any individual consumer, especially one who was unlucky in Y draws or initial X, can go for years without running into a useful interlocutor (one who has been at least as unlucky). One way to ameliorate this problem is to suppose that each consumer communicates with several others each period, so that there is a greater chance of encountering someone at least as lucky as oneself. The right-hand panel of the table presents results when consumers are assumed to interact with 12 other consumers per period (corresponding to monthly interactions if the model period is a year).

As expected, convergence is greatly speeded up in this case: for the $n = 20$ case the average sacrifice value after 100 periods of simulation is half of its value in the baseline model. However, for simulation lengths greater than 500 periods we find that the high-interaction model performs *worse* than the one with fewer interactions. Note that the nature of the bias is that X^* is too low. This presumably reflects the fact that the downward biases in X^* identified in our discussion above are much more problematic in the case with much

Population Statistics by Periods of Simulation Time

	One Interaction/Period				12 Interactions/Period		
	$n = 10$	$n = 20$	$n = 50$		$n = 10$	$n = 20$	$n = 50$
Period 1							
Avg γ:	0.525	0.525	0.525		0.525	0.525	0.525
Avg X^*:	1.950	1.950	1.950		1.950	1.950	1.950
Avg Sacrifice:	1.059	1.059	1.059		1.059	1.059	1.059
Remaining Rules:	1.000	1.000	1.000		1.000	1.000	1.000
Period 100							
Avg γ:	0.268	0.274	0.346		0.263	0.231	0.295
Avg X^*:	1.259	1.315	1.541		1.129	1.213	1.404
Avg Sacrifice:	0.107	0.132	0.308		0.076	0.067	0.186
Remaining Rules:	0.442	0.495	0.865		0.152	0.255	0.637
Period 500							
Avg γ:	0.221	0.220	0.225		0.250	0.224	0.212
Avg X^*:	1.169	1.189	1.222		1.108	1.146	1.192
Avg Sacrifice:	0.042	0.039	0.049		0.060	0.039	0.034
Remaining Rules:	0.125	0.137	0.198		0.058	0.070	0.135
Period 2000							
Avg γ:	0.226	0.237	0.228		0.246	0.233	0.223
Avg X^*:	1.172	1.197	1.198		1.099	1.152	1.179
Avg Sacrifice:	0.024	0.019	0.019		0.061	0.029	0.019
Remaining Rules:	0.045	0.050	0.050		0.027	0.027	0.032
Period 10000							
Avg γ:	0.227	0.228	0.216		N/A†	N/A†	N/A†
Avg X^*:	1.174	1.196	1.215		N/A	N/A	N/A
Avg Sacrifice:	0.020	0.013	0.013		N/A	N/A	N/A
Remaining Rules:	0.018	0.020	0.020		N/A	N/A	N/A

n is the number of periods before social interaction is allowed.

† Because of computer speed limitations, we were unable to simulate 10,000 periods in the multiple-interactions case.

Table 2: Social Learning Simulation Results

more social contact. In particular, rules belonging to imprudent individuals who happen to be lucky in drawing a smooth income stream have a much greater ability to spread widely in the population when such lucky individuals have the chance to encounter a larger number of other consumers. Still, the magnitude of the bias is relatively small: after 2000 periods for the $n = 20$ case, the average sacrifice value is only 0.01 higher in the high-interaction case.[17]

6 What Next?

The work described thus far in this proposal is very preliminary and needs to be extended in a variety of ways.

A simple first step will be to examine the sensitivity of the model's results to variation in model parameters like the time preference rate, the coefficient of relative risk aversion, and the characterization of uncertainty. Particularly interesting will be the relationship between the degree of uncertainty and the speed of convergence. Preliminary exploration suggests that convergence is *much* faster if uncertainty is minimal.

An appealing next step would seem to be to reparameterize the model so that the model's 'period' corresponds to something less than a year. The natural choice would be for the period to correspond to paycheck frequency (24 times per year) or a monthly frequency. Unfortunately, however, at frequencies greater than a year our simple characterization of the process for labor income uncertainty (which roughly matches household-level fluctuations in annual income calculated from the PSID) is probably inappropriate. The most important source of transitory income uncertainty for working consumers is probably the risk of an unemployment spell. Unemployment spells typically last 3-6 months, with typically about 50 percent replacement of earnings via unemployment insurance. The best way to model unemployment spells is probably via a set of unemployment/reemployment hazards, but this introduces another state variable into the problem (employment status), and optimal values for γ and X^* will certainly be different for employed versus unemployed consumers. Thus, moving to a paycheck or monthly frequency introduces considerable complication, along with two more parameters for consumers to learn about; while this extension is probably worth pursuing, is not entirely clear that convergence will be faster (measured in years), given the extra parameters that consumers need to learn.

Another potentially interesting idea is to investigate whether the presence of 'experts' can significantly speed or alter the convergence process. In the Introduction, we mentioned that personal finance books often recommend that people maintain a buffer stock of liquid assets to draw upon in emergencies, and one could interpret the personal finance book as reflecting the experience of agents defined as personal finance 'experts' who seek to have much more than the average amount of social communication with others on the subject of

[17] We have been unable to explore the effects of simulating a larger number of periods for the multiple-interactions case because the computer facilities at our disposal are too slow. For any given value of n, 10000 periods of simulation would take nearly a week.

consumption/saving behavior. It seems possible that if we were to introduce a small number of such highly social agents into the population their presence might (or might not) greatly speed convergence. Eventually one could even imagine these agents charging a fee for their services, though the modelling complications of setting up such a market might prove to be greater than the benefits.

Introducing improvements in the intelligence of our (currently very stupid) artificial agents is inherently appealing, and would almost certainly substantially speed convergence times. One particularly simple idea is to suppose that consumers can engage in a limited form of 'reflection' by which they are able to assess how their current-period utility would have differed if the current-period income shock had taken on any of its possible alternative values.[18] Another idea is to endow consumers with the ability to perform simple counterfactual calculations, so that they can, for example, estimate roughly how much higher their utility might have been had they earned an average income stream as high as that of a luckier interlocutor. In this case we could relax the assumption that all socially communicated information from luckier consumers must be discarded.

Perhaps the most glaring defect in the intelligence of our agents, however, is their almost total lack of memory, either of their own past experience before adopting their current consumption rule, or of the information other agents have communicated to them over time. The most attractive way to remedy this defect would be to build a simple neural network which each consumer could use to keep track both of his own experiences and of any wisdom gleaned from others. We have begun preliminary efforts to construct such neural nets but have concluded that the computational resources required to give each consumer his own individual memory are beyond the capabilities of the computer resources currently available to us.

Given the strong resemblance between the evolution of the population of rules in our model and biological evolution as modelled using genetic algorithms, another avenue we intend to explore is the relationship between our results and comparable work in the literature on genetic algorithms. Our model lacks two of the forces that propel the evolution of genetic algorithms: mutation and genetic recombination. Furthermore, in contrast with the typical approach with genetic algorithms, our model begins with a population that is evenly distributed over the entire 'fitness landscape' defined by γ and X^*. Despite these differences, the progress toward convergence in our model can be interpreted very naturally in terms of the differential 'reproductive success' of the various consumption rules in the population, with social contacts between individuals playing the role of reproductive opportunities for rules. One interesting question to ask is how convergence speed would change if we were to add plausible forms of 'mutation' (spontaneous changes in the rules individuals use) and 'genetic recombination' (rather than adopting someone else's rule entirely, an agent might adopt a rule intermediate between his own and that of his communicant, or might, say, adopt

[18]The term 'reflection' comes from the literature on genetic algorithms in which evolving agents are sometimes allowed to engage in precisely this kind of behavior.

the other person's X^* but not γ.) We could also investigate how robust our findings are to alternative initial distributions of consumers across the fitness landscape. Incorporating some kind of mutation or randomness in the model would clearly be important in addressing the issue of how the population responds to changes in the environment that change the optimal rule. In the extreme case, if there were no mutation or randomness of rules and the entire population had converged to the rule that was optimal in one environment, our current setup would provide no way for the population to migrate to the new rule that is optimal in the new environment.

When we have identified a social learning framework with which we are more fully satisfied, we intend to explore the implications of that framework for empirical work on consumption behavior. A particularly interesting set of questions is the extent to which an individual agent's behavior should be influenced by the recent negative experiences of people with whom the agent has been in social communication. We hope to be able to formulate hypotheses which can be tested with data from the PSID on the assumption that individuals are in social communication with their relatives (siblings interactions will be particularly interesting because the conceptual difficulties posed by bequests may make parent/child interactions difficult to interpret).

7 Future Projects

In part, the exploration of the optimal consumption problem outlined thus far in this proposal is meant as an example of a more general research methodology that we intend to apply in a variety of other contexts.

A first followup project will be to apply our methodology to study consumers' portfolio allocation choices. Our intuition is that optimal portfolio allocation will prove to be even more difficult to divine through personal experience than optimal consumption choice because the uncertainty in stock returns is even greater than that in earnings, and hence social learning is likely to be even more important in this context. While the idea that some investors may base expectations of future returns on past performance is far from new, to our knowledge there has been very little investigation of these issues using social learning/genetic algorithm frameworks. A recent paper by Arthur et. al. (1997) does investigate the behavior of a population of forecasting rules evolving via a genetic algorithm, but that paper does not investigate the question of individuals' portfolio allocations between a safe and a risky asset, which is the heart of the microeconomic manifestation of the equity premium puzzle, and which would be the principal focus of our efforts.

More broadly, there are many spheres in which rational responses of economic agents to changes in the environment appear to have emerged, but only with extremely long lags (Lindbeck (1995); Lindbeck, Nyberg, and Weibull(1997); and references therein for examples). While these authors have explained the delays as the result of a gradual erosion of social norms, social learning seems to be a plausible alternative explanation. Our intended first application of the social learning framework to such an issue will be with respect to the evolution of the personal bankruptcy rate in the United States. In a series of recent papers, Michelle White (1998) has shown that current US bankruptcy laws are so generous that approximately 15 percent of US households could immediately increase their net worth by declaring bankruptcy, and almost 50 percent could benefit from bankruptcy if they first engaged in a modest amount of bankruptcy planning. From this perspective, the mystery is not why the personal bankruptcy rate has been increasing at an annual rate of 30 percent over the last few years despite magnificent economic conditions, but rather why the bankruptcy rate remained so low so long. Bankruptcy lawyers' lobbying groups have argued that the generosity of the bankruptcy law has nothing to do with the increase in bankruptcy rates, because the law has remained largely unchanged since 1978, when the current debtor-friendly law replaced a much more punitive bankruptcy system. While the *uptrend* in bankruptcy rates appears to have begun shortly after the change in bankruptcy law (dating the beginning of the uptrend is a bit difficult because of interference from the 1980-82 recessions), it is true that the standard economic theory of bankruptcy implies that there should have been an immediate avalanche of bankruptcies in 1979 rather than a gradual uptrend. White and others who believe that the fundamental cause of high bankruptcies is the generosity of the bankruptcy law have argued that the delay is attributable to the gradual erosion of an ill-defined 'social stigma' associated with bankruptcy. An alternative explanation is that the evolution of bankruptcy rates reflects a gradual process of social learning about the increased generosity of the bankruptcy law, and this a natural next application of our framework.

The scope for other applications of frameworks like the one described in this proposal is vast, even outside of the game theory literature where similar ideas are already gaining ground. We intend to encourage further development of such work by making the computer software used to solve and simulate our model publicly available when the projects are done, and available on request even before then. In anticipation of such future uses, we have tried to keep the structure our programs as flexible and general-purpose as possible.

References

Arthur, W. B. (1993): "On Designing Economic Agents That Behave Like Human Agents," *Journal of Evolutionary Economics*, 3, 1–22.

Arthur, W. B., J. H. Holland, B. LeBaron, R. Palmer, and P. Tayler (1997): "Asset Pricing Under Endogenous Expectations in an Artificial Stock Market," in *The Economy as an Evolving Complex System II*, ed. by W. B. Arthur, S. N. Durlauf, and D. Lane. Addison Wesley Longman.

Binmore, K., and L. Samuelson (1997): "Muddling Through: Noisy Equilibrium Selection," *Journal of Economic Theory*, 74, 235–65.

Browning, Martin J., and Annamaria Lusardi (1996): "Household Saving: Micro Theories and Micro Facts," *Journal of Economic Literature*.

Carroll, C. D. (1992): "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity*, 1992(2), 61–156.

——— (1994): "How Does Future Income Affect Current Consumption?," *The Quarterly Journal of Economics*, CIX(1).

——— (1996): "Buffer Stock Saving: Some Theory," *Manuscript, Johns Hopkins University, January.*

——— (1997a): "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics*, CXII(1), 1–56.

——— (1997b): "Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second Order Approximation)," *Manuscript, Johns Hopkins University.*

Carroll, C. D., and M. S. Kimball (1996): "On the Concavity of the Consumption Function," *Econometrica*, 64(4), 981–992.

Deaton, A. S. (1991): "Saving and Liquidity Constraints," *Econometrica*, 59, 1221–1248.

——— (1992): *Understanding Consumption*. Oxford University Press, New York.

Friedman, M. A. (1957): *A Theory of the Consumption Function*. Princeton University Press.

Hall, R. E. (1978): "Stochastic Implications of the Life-Cycle/Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*, 96, 971–87.

Hall, R. E., and F. Mishkin (1982): "The Sensitivity of Consumption to Transitory Income: Evidence from PSID Households," *Econometrica*, L, 461–81.

HUBBARD, R. G., J. S. SKINNER, AND S. P. ZELDES (1994): "The Importance of Precautionary Motives for Explaining Individual and Aggregate Saving," in *The Carnegie-Rochester Conference Series on Public Policy*, ed. by A. H. Meltzer, and C. I. Plosser, vol. 40, pp. 59–126.

——— (1995): "Precautionary Saving and Social Insurance," *Journal of Political Economy*, 103, 330–399.

LINDBECK, A. (1995): "Welfare State Disincentives with Endogenous Habits and Norms," *Scandinavian Journal of Economics*, 97(4), 477–94.

LINDBECK, A., S. NYBERG, AND J.-W. WEIBULL (1997): "Social Norms and Economic Incentives in the Welfare State," *Industrial Institute for Economic and Social Research.*

MARCET, A., AND T. J. SARGENT (1991): "Convergence of Least Squares Learning Mechanisms in Self Referential Linear Stochastic Models," *Journal of Economic Theory*, 48, 337–68.

ROTH, A., AND I. EREV (1995): "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," *Games and Economic Behavior*, 8, 164–212.

SAMUELSON, L. (1997): *Evolutionary Games and Equilibrium Selection.* MIT Press, Cambridge.

SARGENT, T. J. (1993): *Bounded Rationality in Macroeconomics.* Oxford University Press, Oxford.

SHILLER, R. J. (1997): "Human Behavior and the Efficiency of the Financial System," *Manuscript, Yale University.*

SOULELES, N. S. (1995): "The Response of Household Consumption to Income Tax Refunds," *Manuscript, the Massachusetts Institute of Technology.*

THALER, R. (1994): "Psychology and Savings Policies," *American Economic Review.*

WEIBULL, J. (1995): *Evolutionary Game Theory.* MIT Press, New York.

WHITE, M. (1998): "Why It Pays to File For Bankruptcy: A Critical Look at Incentives Under U.S. Bankruptcy Laws and a Proposal for Change," *University of Michigan Working Paper*, 98-2.

YOUNG, H. P. (1997): *Learning and Evolution in Games.* Princeton University Press.