ELSEVIER

JOURNAL OF
Economic
Dynamics
& Control

Rules of thumb in macroeconomic equilibrium
A quantitative analysis

Per Krusell[a,b], Anthony A. Smith, Jr.[*,c]

[a]*University of Rochester, Rochester NY 14627, USA*
[b]*Institute for International Economic Studies, S-10691 Stockholm, Sweden*
[c]*Carnegie Mellon University, Pittsburgh, PA 15213, USA*

Abstract

We investigate the robustness of the general equilibrium stochastic growth model to the introduction of small costs of behavioral sophistication. Consumers choose amongst savings rules which vary in sophistication and effort cost. In our model, a given consumer's gain from using a sophisticated rule is higher when other consumers use simple rules. Thus, decentralization makes it harder for simple rules to survive than in similar decision-theoretic models. Nevertheless, we find: (i) that sophisticated behavior, which we model as fully unrestricted, occurs in equilibrium only if its relative effort cost is extremely low; and (ii) that rule-of-thumb economies can generate aggregate time series that differ substantially from those of the standard model.

Key words: Stochastic growth model; Robustness; Rules of thumb; Decentralized decision-making
JEL classification: E00; E32

1. Introduction

Many recent macroeconomic theories emphasize the intertemporal nature of the decisions faced by economic agents. The behavioral predictions of these

theories are often quite sophisticated. Motivated by the perception that the behavior of actual economic agents often seems less sophisticated than required by these theories, some researchers have introduced various forms of 'bounded rationality' into macroeconomic modeling. One example of this line of research is the literature on learning (see, e.g., Evans and Honkapohja, 1993; Sargent, 1992).

In this paper, we analyze a dynamic equilibrium environment which is similar to the workhorse stochastic growth model, except in one important respect: agents face explicit costs of following sophisticated behavioral rules. Given that we think that real-world agents do face various forms of costs – costs of computation, costs of acquiring information, costs of implementing decisions, etc. – we view this type of exercise as an important test of the *robustness* of the standard model as a framework for studying macroeconomic phenomena. The purpose of the paper is thus to analyze how the introduction of these costs into a decentralized version of the stochastic growth model affects the model's predictions for individual as well as for aggregate variables. A central aspect of this robustness exercise is that we consider a decentralized economy. The purpose of this decentralization is to allow any given agent to respond to, and possibly take advantage of, the unsophisticated behavior of other agents.

We find that extremely small relative costs of unrestricted behavior, i.e., the behavior that would be optimal for the consumers in the absence of any costs, lead to equilibrium outcomes dominated by unsophisticated – 'rule-of-thumb' – behavior. By extremely small we mean costs that are less than a tenth of a percent of per period consumption. For example, if the rule of thumb is an (optimally chosen) savings rate that is required to be constant over time, then all consumers use this rule of thumb so long as its relative cost advantage exceeds as little as 0.03% of per period consumption in a reasonably parameterized model economy.

At the same time, rules of thumb can give rise to aggregate time series behavior that is quite different in several dimensions than that arising in an economy without decision rule costs. For example, one of the main successes of stochastic intertemporal optimization has been to explain why aggregate consumption is considerably less volatile in percentage terms than aggregate income. With even very small costs of decision rule sophistication, this implication may disappear: although the agent would, ideally, like to smooth consumption relative to income, the utility gain from smoothing is so small that many much cruder rules simply outperform consumption smoothing on a net-of-costs basis. In our examples, we show this using the constant-savings-rate rule and a constant-capital-stock rule: the former makes consumption, investment, and income equally volatile in percentage terms, while the latter leads to perfectly smooth investment, making consumption more volatile than income.

We also find, as a general point, that the aggregate time series behavior of the model is very sensitive to the precise set of rules available to the agents and to

the relative costs of these different rules. In particular, equilibrium outcomes in which rule-of-thumb behavior dominates need not display time series behavior which differs substantially from that displayed by an economy without decision rule costs. For example, the stochastic growth model has the property that fully unrestricted behavior is very well approximated by making consumption a log-linear function of the agent's state variables (individual and aggregate capital and the aggregate productivity shock). Hence, with such a rule, the aggregate time series would be very similar to those of the standard model. In this paper we make an even stronger point: there exist rules which are much simpler than the log-linear rule, and which reproduce the broad time series features of the standard model, including the consumption smoothing property. These simple rules have only one parameter, and they include what traditionally has been called a partial-adjustment rule for capital, as well as a rule setting consumption equal to a constant times the agent's current holdings of capital.

The economic model we consider is the stochastic neoclassical growth environment populated with a large number of consumers who are identical as of the beginning of time. These agents receive labor and capital income from a perfectly competitive production sector. Since we assume that there is no labor–leisure choice, the only decision for the agent is intertemporal: how much to consume and how much to save at each point in time. In a decentralized context, this decision requires, formally, an analysis of a stochastic optimal control/dynamic programming problem (see Brock and Mirman, 1972; Stokey and Lucas, 1989). Its solution takes the form of a time-independent decision rule which maps the current level of capital holdings by the agent, the aggregate level of capital, and the aggregate productivity shock into a choice for current consumption and savings. In this paper we assume that this decision rule is 'hard to implement' in this fully unrestricted form: in particular, there are more restricted types of behavior which require less effort. These restrictions take the form of rules of thumb, and they include time-and-state-independent savings rates or constant capital stocks, restrictions to linear functional forms and to limited dependence on the elements of the state vector, and so on.

An equilibrium in our economy is then one where the agent chooses among the set of available decision rules, weighing the effort costs of an unrestricted behavioral rule against the utility losses implied by the rules of thumb. In an environment in which agents choose between the two types of decision rules, such as, for example, a sophisticated (fully unrestricted) rule and an unsophisticated rule of thumb characterized by a constant savings rate, three types of equilibria can emerge. If the relative effort cost associated with using the sophisticated rule, η, is smaller than some number $\underline{\eta}$, then all agents use the sophisticated rule. If the cost is larger than some other, higher number, $\bar{\eta}$, then all agents use the simple rule. Finally, if $\eta \in (\underline{\eta}, \bar{\eta})$, then a fraction $g(\eta)$ of the agents use the sophisticated rule and the rest use the rule of thumb. We find g to be a continuous and monotonically decreasing function. There is, hence, the

possibility that agents sort themselves ex ante into sophisticated and less sophisticated types. In this kind of equilibrium, the sophisticated agents are able to take advantage of their better ability to adjust consumption and investment behavior in response to shocks and changes in the return to capital, but these advantages are just enough to equal the costs borne by using the complicated rule.

Our results are related to earlier findings by Lucas (1987), Cochrane (1989), and Smith (1990, 1992). Lucas, first, compared the utility of a completely flat consumption profile with that of a fluctuating profile, with the magnitude of the fluctuations chosen to match that of observed aggregate consumption series. He found the difference in utility to be very small. Cochrane and Smith then looked at a larger class of rules of thumb, and argued more explicitly in terms of the relative costs of using the different rules: even if the relative cost of following fully flexible decision rules rather than simpler rules of thumb is very small, then in economies where all agents use the same decision rule, we would expect the agents to become rule-of-thumb decision makers.

Although quite suggestive, these findings clearly have a problem: they are not derived in decentralized economies.[1] In this context, this shortcoming is potentially very troublesome: although it may not pay off for all agents to switch to sophisticated rules in a coordinated fashion, a given small agent operating in an economy inhabited by rule-of-thumb decision makers may be able to obtain much higher payoffs from switching to a sophisticated rule by taking advantage of the simple rules used by other agents. In other words, the one-agent experiments performed by Lucas, Cochrane, and Smith suggest a lack of robustness of the model which may disappear once one recognizes that sophisticated agents can 'exploit' the suboptimal behavior of others.

Our results here show that, qualitatively, this mechanism is indeed operating: the cutoff costs above which we would observe only rule-of-thumb decision makers in the decentralized economy are one-and-a-half times larger (in welfare terms) than the welfare gains that obtain in a centralized setting such as that explored by Smith (1990, 1992). The amount of the change, however, is not large enough to shake the main implications of this line of work regarding the robustness of the stochastic intertemporal optimization framework.

Our approach to modelling lack of full consumer sophistication is not to give up rationality and the consumer's understanding of the world in which he lives, but instead to recognize the existence of various costs associated with making and implementing decisions. This view is consistent with 'rule-rationality' as discussed by Robert Aumann in his 'Perspectives on Bounded Rationality' lecture (Aumann, 1992), and it has also been explored independently in the context of

---

[1] More precisely, they are decision-theoretic or, alternatively, they assume that all agents must behave in the same way.

a more standard game-theoretic setting (see Rosenthal, 1991). In contrast, the learning literature postulates that the agent is not fully rational: he uses, say, a forecasting device which is postulated in an ad hoc fashion, and which could be improved upon by an agent who forecasts in an unrestricted way and understands the behavior of the rest of the agents in the economy.

What is ad hoc in this paper, and therefore in its current form a weakness, is the structure we impose for the costs associated with the different decision rules. In particular, the complexity of the rules is not made formal by appeal to numerical procedures involved in complex computations (as in, for example, Hurwicz, Reiter, and Saari, 1992). Ideally, one would like the costs to be more easily identifiable in order to impose further empirical discipline on the analysis – as we formulate the costs now, only introspection can tell whether a given cost is empirically plausible. This weakness is especially important here, since we find that the model's aggregate behavior depends critically on the precise set of rules available to consumers, together with the rules' associated costs. For example, if the constant-savings-rate rule is the cheapest possible rule, the aggregate time series will look quite different than if the partial-adjustment rule is the cheapest rule.

Moreover, we also do not attempt to solve the fundamental infinite-regress problem associated with 'deciding how to decide how to ...' (see Lipman, 1991). Indeed, one interpretation of our costs is simply that they are associated with the implementation of complex behavior, and not with evaluating its benefits. Our main goal in this paper, however, is a preliminary quantitative evaluation of the macroeconomic effects of allowing rule-of-thumb behavior. For this purpose, we prefer our simple operationalization of behavioral costs, reserving a search for better structural formulations for future research. It should also be pointed out that our structure does provide more structure than do models with purely 'irrational' behavior taken as exogenous inputs into the analysis.[2]

An additional result of our analysis is that the kind of decentralized decision rule selection we consider gives rise to a kind of 'externality'. The reason why social benefits from choosing decision rules do not coincide with social costs is that the consumption possibility set of the less sophisticated agents depends on aggregate choices. For example, the set of attainable levels of investment and consumption for an agent using a constant-savings-rate rule depends on the process for his income, which in turn depends on aggregate choices. As a result, the competitive equilibrium fraction g is not, in general, 'socially optimal'. More

---

[2] For examples see Akerlof and Yellen (1985a, b), Caballero (1991), Campbell and Mankiw (1990), Cochrane (1989), De Long et al. (1990), Haltiwanger and Waldman (1985, 1991), Ingram (1990), Mirrlees and Stern (1972), Smith (1990, 1992), and Weil (1991). In Evans and Ramey (1992) and Crettez and Michel (1990), agents are allowed to choose the level of sophistication of their actions, but in economic environments that differ substantially from the stochastic neoclassical growth environment studied here.

precisely, if the relative effort cost $\eta \in (\underline{\eta}, \bar{\eta})$, then we find that an ex ante social welfare criterion would lead to a fraction of sophisticated agents that is smaller than that resulting in competitive equilibrium. It is also true that a move to the fraction $g^* < g$ that maximizes ex ante utility makes both types of agents better off ex post. The deviation from a socially optimal fraction of sophisticated agents is quantitatively small and, in economic terms, has only a small effect on agents' welfare.[3]

We describe the setup with decision rule selection in Section 2, and in Section 3 we describe our results. Finally, Section 4 draws tentative conclusions and makes some suggestions for future research.

2. The model

In this section we describe the economy with decision rule selection. To simplify the notation in our presentation, we focus on a simple case: one in which each agent has a choice between fully unrestricted behavior and a constant-savings-rate rule (with the savings rate chosen optimally). It is straightforward to modify this framework to handle other pairs of decision rules, and in Section 3 we provide results for economies with other pairs of decision rules as well.

2.1. The economic environment

We assume that there is a continuum, say of measure 1, of identical agents. Agents derive utility from consumption goods of which there is one type per period. There is a constant-returns-to-scale production technology that transforms capital and labor at time t into consumption good at time t. We assume that this production function is given by

$$Y_t = K_t^\alpha H_t^{1-\alpha} \exp(z_t),$$

where Y_t is aggregate output, K_t is the aggregate capital stock, and H_t is the total amount of labor, which we assume is inelastically supplied at the (total and per-capita) amount of 1. The technology shock z_t follows the stochastic process

---

[3] This result is not necessarily at odds with the finding in Akerlof and Yellen (1985a) that suboptimal (near-rational) behavior in an economy with externalities can lead to a first-order decrease in aggregate (social) welfare. In our framework, the welfare loss associated with the competitive equilibrium may indeed be first-order, but is nonetheless quantitatively small.

$z_{t+1} = \rho z_t + \varepsilon_{t+1}$, with $\varepsilon_{t+1} \sim \text{iidN}(0, \sigma_\varepsilon^2)$. The parameters α and ρ satisfy $\alpha \in (0, 1)$ and $\rho \in (-1, 1)$.

Intertemporal production (i.e., capital accumulation) takes place according to

$$K_{t+1} = (1 - \delta)K_t + X_t,$$

where X_t is total investment, K_0 is given, and $\delta \in (0, 1)$. The resource constraint in the economy reads

$$C_t + X_t \leqslant Y_t,$$

where C_t is total consumption.

We define the following functions:

$$MPK_t(K_t, z_t) = \alpha K_t^{\alpha-1} \exp(z_t), \qquad MPL_t(K_t, z_t) = (1 - \alpha)K_t^{\alpha} \exp(z_t).$$

These variables, the marginal productivity of capital and labor, respectively, are needed in the specification of the agent's consumption possibility set, C. This consumption possibility set contains the admissible processes of consumption, c, and effort levels, $e \in \{0, 1\}$. The description of preferences that will be made below shows how effort affects utility. The general idea here is that effort can be expended in return for a less restrictive set of possible consumption processes. There is a fundamental difference between the consumption possibility set used here and those of standard, Arrow–Debreu economies. In the latter, the consumption possibility set is simply an appropriately chosen space that involves nonnegativity and certain more technical conditions needed for the existence of a price functional that has an inner product representation. The consumption possibility set here, on the other hand, depends on aggregate technology variables, and in its specification it uses an individual version of the intertemporal technology.

The set C depends on the stochastic processes MPK_t and MPL_t, since the less sophisticated consumption rules express consumption as a function of these variables (later to be interpreted as 'income'). The set is defined as follows:[4]

$$C = \{(c, e) \in \mathscr{C} \times \{0, 1\} : \text{either } e = 1 \text{ and } c \in \mathscr{C} \text{ or } e = 0 \text{ and } c \in \mathscr{C}_u\},$$

---

[4] We omit the measurability restrictions on the stochastic processes described. These restrictions are straightforward but tedious to state, since they involve the processes MPK_t and MPL_t.

where $\mathscr{C}$ is a suitable set of stochastic processes (measurable and nonnegative) and where

$$\mathscr{C}_u = \{c \in \mathscr{C}: \exists s \in [0, 1] \text{ and } (x, k) \in \mathscr{C}^2 \text{ such that}$$

$$c_t + x_t = MPK_t(K_t, z_t)k_t + MPL_t(K_t, z_t),$$

$$k_{t+1} = (1 - \delta)k_t + x_t \text{ with } k_0 = K_0, \text{ and}$$

$$c_t/x_t = 1/s - 1\}.$$

In the definition of $\mathscr{C}_u$, the lower-case variables c_t, x_t, and k_t denote, respectively, period t values for consumption, investment, and the capital stock for an individual agent. In other words, if effort is expended ($e = 1$), the consumption possibility set is unrestricted, whereas in the case no effort is expended ($e = 0$), the agent is restricted to consumption paths such that the ratio of consumption to total 'private resources' at every point in time is equal to a number $1 - s$ that is constant across time and states but can be chosen freely. The private resources are selected as the sum of the products between each factor and its marginal product, using the capital path that is implied by the consumption/savings behavior given by s and using an initial value of capital that satisfies $k_0 = K_0$. In Section 3, we will consider other kinds of simple rules as well, and it should be clear to the reader how $\mathscr{C}_u$ has to be formulated in each case.

Finally, the agent's preferences over consumption streams and effort are represented by

$$E_0 \sum_{t=0}^{\infty} \beta^t u(c_t) - \eta e,$$

where η is the utility loss of a unit of effort and E_0 is the expectation operator conditional on the time 0 information set.

2.2. Equilibrium

We use a competitive equilibrium concept. The final output in period t is produced in firms solving

$$\max_{K_t, H_t} Y_t - r_t K_t - w_t H_t$$

subject to

$$Y_t = K_t^\alpha H_t^{1-\alpha} \exp(z_t),$$

where r_t and w_t are the relative prices of capital and labor services. Because of the constant returns assumption, these will have to equal MPK_t and MPL_t, respectively, in equilibrium. Consumers solve

$$\max_{c, e, k, x, y} E_0 \sum_{t=0}^{\infty} \beta^t u(c_t) - \eta e$$

subject to

$$c_t + x_t = r_t k_t + w_t \equiv y_t,$$

$$k_{t+1} = (1 - \delta)k_t + x_t,$$

$$k_0 = K_0,$$

$$(c, e) \in C.$$

It is hence implicit in this problem from the last of the constraints that the agent has as an option to select $e = 0$ and a corresponding constant savings fraction s. Also note that, in the case $e = 0$, the capital accumulation path implicit in the set C will be consistent with the paths of the explicit constraints above, since in equilibrium rental rates coincide with marginal products.

Note that the sophisticated agent, i.e., an agent who selects $e = 1$, does not have access to complete markets. This is not, however, a restriction: if a sophisticated agent were to trade in some contingent claim not available through the present market structure, he would have to trade with another sophisticated agent. This observation follows from our assumption that agents who choose $e = 0$ have access only to their special decision rule: allowing such an agent to trade in additional claims would violate this assumption. But since sophisticated agents are all alike, and since their preferences are convex, they will prefer to be the same and hence not need a more elaborate market structure.

Observe that the consumer's problem can be solved in two steps. In the first step, the consumer selects consumption paths corresponding to each of the cases $e = 0$ and $e = 1$ separately. In the second step, the consumer chooses the value of e that is associated with the highest utility as calculated in the first step. The first step involves solving the sophisticated agent's problem (PS):

$$v_S \equiv \max_{c, k, x, y} E_0 \sum_{t=0}^{\infty} \beta^t u(c_t) - \eta$$

subject to

$$c_t + x_t = r_t k_t + w_t \equiv y_t,$$

$$k_{t+1} = (1 - \delta)k_t + x_t,$$

$$k_0 = K_0,$$

$$c \in \mathscr{C},$$

and solving the unsophisticated agent's problem (PU):

$$v_U \equiv \max_{c,\,s,\,k,\,x,\,y} \mathrm{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t)$$

subject to

$$c_t + x_t = r_t k_t + w_t \equiv y_t,$$

$$k_{t+1} = (1 - \delta)k_t + x_t,$$

$$k_0 = K_0,$$

$$c \in \mathscr{C} \quad \text{with} \quad c_t/y_t = 1 - s.$$

These problems give rise to laws of motion for the total capital stock of sophisticated agents, K_{St}, and of unsophisticated agents, K_{Ut}. The aggregate capital stock is $K_t = \theta K_{St} + (1 - \theta)K_{Ut}$, θ being the fraction of sophisticated agents in equilibrium.

The sophisticated consumer's problem is recursive. This follows, first, since the prices are functions of K_t and z_t. Second, these variables are composed of K_{St}, K_{Ut}, and z_t, which jointly follow a recursive law of motion. The variable z_t does so by assumption, K_{St} does so by hypothesis, and K_{Ut+1} equals

$$(1 - \delta)K_{Ut} + s[MPK_t(\theta K_{St} + (1 - \theta)K_{Ut}, z_t)K_{Ut}$$

$$+ MPL_t(\theta K_{St} + (1 - \theta)K_{Ut}, z_t)],$$

which is an expression in K_{St}, K_{Ut}, and z_t only. Because of the recursive, standard nature of this problem, it generates a unique solution in the form of a decision rule for capital f_S, with $k_{t+1} = f_S(k_t, K_{St}, K_{Ut}, z_t)$.

The relevant state variable of the economy is the triplet (K_{St}, K_{Ut}, z_t). The aggregate laws of motion of the two capital stocks are denoted F_S and F_U:

$$K_{St+1} = F_S(K_{St}, K_{Ut}, z_t), \qquad K_{Ut+1} = F_U(K_{St}, K_{Ut}, z_t).$$

As argued above, the rental rates are both functions of the aggregate state.

We will focus on equilibria in which all unsophisticated agents choose the same value of s – this is motivated by our numerical finding that the unsophisticated agent's problem has a unique solution. We denote the unsophisticated agent's decision rule f_U, with

$$k_{t+1} = f_U(k_t, K_{St}, K_{Ut}, z_t)$$

$$= (1 - \delta)k_t + s[MPK_t(\theta K_{St} + (1 - \theta)K_{Ut}, z_t)k_t$$

$$+ MPL_t(\theta K_{St} + (1 - \theta)K_{Ut}, z_t)].$$

Our equilibrium can hence be defined as follows:

Definition 1.   A competitive equilibrium is a pair of pricing functions $r(K_S, K_U, z)$ and $w(K_S, K_U, z)$, a savings rate s, a fraction θ, aggregate laws of motion $F_S(K_S, K_U, z)$ and $F_U(K_S, K_U, z)$, and decision rules $f_S(k, K_S, K_U, z)$ and $f_U(k, K_S, K_U, z)$, such that

(1) *Sophisticated agents maximize: f_S solves (PS).*

(2) *Unsophisticated agents maximize: s and f_U solve (PU).*

(3) *Firms maximize: $r = MPK$ and $w = MPL$.*

(4) *Consistency holds: $F_S(K_S, K_U, z) = f_S(K_S, K_S, K_U, z)$ and $F_U(K_S, K_U, z) = f_U(K_U, K_S, K_U, z)$.*

(5) *Agents make optimal decision rule choices: (a) $\theta \in (0, 1) \Rightarrow v_S = v_U$, (b) $v_S > v_U \Rightarrow \theta = 1$, (c) $v_S < v_U \Rightarrow \theta = 0$.*

It should be noted, again, that although this definition is written specifically for the case of a constant-savings-rate rule, it is straightforward to amend it to allow for different, and possibly more, types of rules.

2.3. Some qualitative equilibrium characteristics

We have not formally proved existence of the equilibrium as defined. However, an algorithm for computing the equilibrium numerically converges

successfully for a variety of starting values, suggesting that equilibria do exist and are unique. The methods for computing equilibria are described in the Appendix.

In this section, we describe some of the qualitative characteristics of the computed equilibria. These qualitative characteristics apply not only to the economy in which consumers choose between an unrestricted rule and a constant-savings-rate rule, but also to economies with other pairs of decision rules (in particular, those for which we report numerical results in Section 3).

In all of the numerical examples that we have computed, we find that the fraction θ of sophisticated agents is a weakly decreasing function $g(\eta)$ of the effort cost. We also find that there always exists a pair of costs $\bar{\eta} > \underline{\eta} > 0$ such that $g(\eta) = 1$ if $\eta \leqslant \underline{\eta}$, $g(\eta) = 0$ if $\eta \geqslant \bar{\eta}$, and $g(\eta)$ is continuous and strictly decreasing over $[\underline{\eta}, \bar{\eta}]$.

To gain more insight into the determination of θ, it is useful to think about the utility achieved by the two kinds of agents (gross of effort cost) as a function of the equilibrium fraction of sophisticated agents: $U_S(\theta)$ and $U_U(\theta)$. It will always be true that sophisticated agents have higher utility (gross of effort cost), i.e., $U_S(\theta) > U_U(\theta)$ for all θ. We also find that $U_S(\theta)$ is a continuous and monotonically decreasing function, whereas $U_U(\theta)$ is continuous and monotonically increasing. In other words, sophisticated agents are better off with fewer sophisticated agents in the economy, whereas the reverse is true for the unsophisticated agents. There is, hence, a sense in which sophisticated and unsophisticated agents thrive off of each other, and hence 'disagree' over the desirability of changes in θ.[5]

Following these observations, and the equilibrium definition, the cost η that supports the fraction θ is simply $U_S(\theta) - U_U(\theta)$. In particular, the cutoff levels for the effort costs satisfy $\underline{\eta} = U_S(1) - U_U(1)$ and $\bar{\eta} = U_S(0) = U_U(0)$. Since $\eta = U_S(\theta) - U_U(\theta)$ is continuous and monotonically decreasing in θ over $[0, 1]$, the competitive equilibrium fraction θ of sophisticated agents must fall as η rises over $[\underline{\eta}, \bar{\eta}]$.

3. Results

In this section we present our quantitative results. First, we focus on the case of the constant-savings-rate rule used in the exposition in the previous section. We use this rule to make two points. The first of these points is that equilibria

---

[5] In the terminology of Haltiwanger and Waldman (1985, 1991), the economy with decision rule selection is characterized by 'congestion effects' and 'strategic substitutability'. Haltiwanger and Waldman argue that in such environments sophisticated agents have a disproportionately large impact on the economy's equilibrium. Our goal in this paper is to measure these effects quantitatively in a specific general equilibrium environment.

will be characterized by rule-of-thumb behavior unless the cost of sophistication is extremely small. The second point we make with this rule is that it does also give rise to aggregate time series statistics which are quite different from those of the standard model. We choose this particular rule to make these points because of its extreme simplicity – it requires knowledge only of present income and it has only one parameter – and because it has strong roots in the traditional macroeconomic literature.[6]

Second, in the next subsection we consider a variety of other simple rules. These other rules serve to show that our key result – that rules of thumb are likely to dominate the equilibrium behavior of agents unless their relative cost advantage is extremely small – is quite general. The rules we choose also show how aggregate time series can look quite different depending on the types of rules considered. The rules which we look at encompass both consumption smoothing and investment smoothing.

Third and finally, this section contains some remarks on properties of optimal allocations. This section is relevant in our economies because of the externality that enters via the agents' consumption possibility sets.

3.1. The constant-savings-rate rule

Our specific functional form for preferences has either $u(c_t) = (1 - v)^{-1} c_t^{1-v}$, where $v > 1$ is the coefficient of relative risk aversion, or $u(c_t) = \log(c_t)$. In the numerical computations we report below, the structural parameter values are set as follows: $\alpha = 0.36$, $\beta = 0.99$, $\delta = 0.025$, $\rho = 0.95$, and $\sigma_\varepsilon = 0.007$. These are standard choices for these parameter values in the economy without decision rule costs (see, for example, Hansen and Wright, 1992). These parameter values reproduce certain observed long-run time-series statistics such as capital's share of income and the capital–output ratio. In addition, in the economy without decision rule costs, these parameter values generate equilibrium time series whose persistence and volatility roughly match those of observed U.S. aggregate time series.[7] The only parameter which is difficult to pin down on the basis of first- and second-moment properties is the coefficient of relative risk aversion v. We therefore report sensitivity analysis with respect to this parameter: we let v take on the values 1, 2, 3, and 5 (with $v = 1$ corresponding to logarithmic preferences). The model is solved for 101 different values of the effort cost parameter η, with implied values of θ ranging from 0 to 1 in increments of 0.01.[8]

---

[6] Many Keynesian models characterize consumer savings behavior using this rule. Moreover, Solow (1956) uses this rule in his seminal paper on neoclassical economic growth.

[7] The economies with decision rule costs reproduce the same long-run statistics but differ in their second-moment properties.

[8] The Appendix shows how the equilibrium is computed by selecting a value of θ and backing out the parameter η that supports the given θ as a competitive equilibrium.
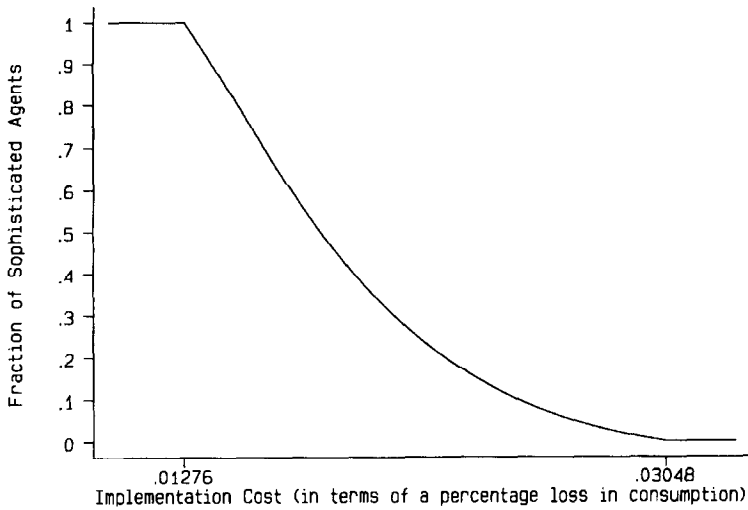
Fig. 1. Competitive equilibrium fraction of sophisticated agents vs. implementation cost.

We make the cost η economically interpretable by expressing the loss in utility (gross of effort cost) in terms of a uniform percentage decrease in consumption across all periods of the planning horizon. That is, we solve for λ in the following equation:

$$E_0 \sum_{t=0}^{\infty} \beta^t u(c_{St}) - \eta = E_0 \sum_{t=0}^{\infty} \beta^t u((1 - \lambda)c_{St}),$$

where c_S is the equilibrium consumption sequence chosen by the sophisticated agent. Note that the term $1 - \lambda$ on the right-hand side of the equation factors out for the utility functions that we use. When $v > 1$, the solution for λ is $\lambda = 1 - (1 - \eta/U_S)^{1/(1-v)}$, where U_S is the equilibrium utility (gross of effort cost) of the sophisticated agent. For logarithmic preferences ($v = 1$), the solution is: $\lambda = 1 - \exp(-\eta(1 - \beta))$. The cost η hence is equivalent to a $100 \times \lambda\%$ decrease in consumption per period, uniformly across all periods and states.

In Fig. 1 (which is based on $v = 2$), we show the equilibrium fraction of sophisticated agents as a function of the effort cost η, expressed in terms of a per period percentage loss in consumption as calculated above. We see that if the cost is less than 0.013% of per period consumption, then all agents use the sophisticated rule in equilibrium. If the cost on the other hand is greater than 0.03%, then all agents use the unsophisticated decision rule. Thus, if the cost is smaller than 0.03% of per period consumption, at least some agents use the sophisticated rule in equilibrium. We use this number as a measure of the
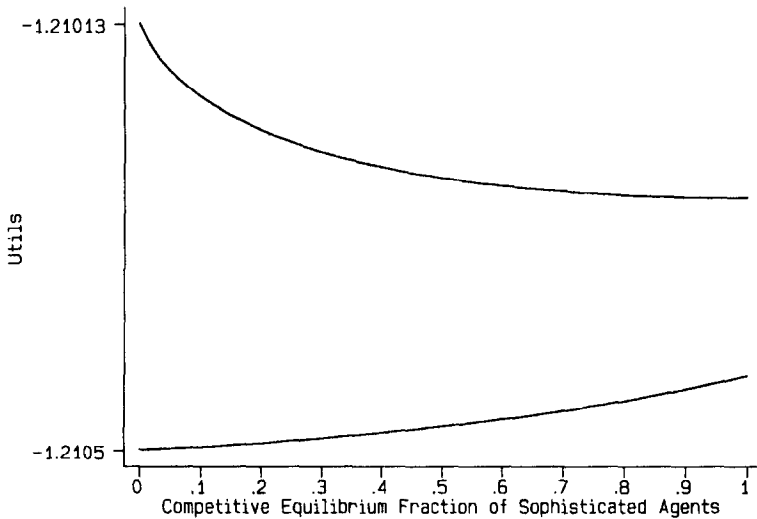
Fig. 2. Utility of sophisticated agents (top) and unsophisticated agents (bottom) in competitive equilibrium.

'incentive to optimize' in this economic environment. This number is very small: assuming annual per capita consumption expenditures of $20,000, no consumers use the sophisticated rule in equilibrium unless its relative cost is less than roughly $6 per year.

It is useful at this point to make a comparison between the results obtained here and the corresponding results in decision-theoretic models. As pointed out earlier, one interpretation of the decision-theoretic setups is that they restrict all agents to using the same rule. The cutoff cost above which those models predict rule-of-thumb behavior then corresponds to the gain a typical agent would realize if all agents were to switch simultaneously to the fully optimal rule. For our present environment, we calculate that this cutoff cost is 0.018% of per period consumption. The cutoff cost in the decentralized model of 0.03% (which is the value above which all agents choose to use the constant-savings-rate rule) is therefore roughly one-and-a-half times larger than the cutoff cost computed under the alternative assumption that agents do not explicitly choose decision rules (and all agents use the same rule).

Fig. 2 graphs the utility (gross of effort cost) for sophisticated and unsophisticated agents for a range of values of the equilibrium fraction of sophisticated agents, θ. The difference between the curves is the utility cost of the effort expense associated with the sophisticated rule that generates the given θ in equilibrium.

We also simulated time series for investment, consumption, and income of the two kinds of agents in the equilibrium with decision rule selection. Fig. 3, which
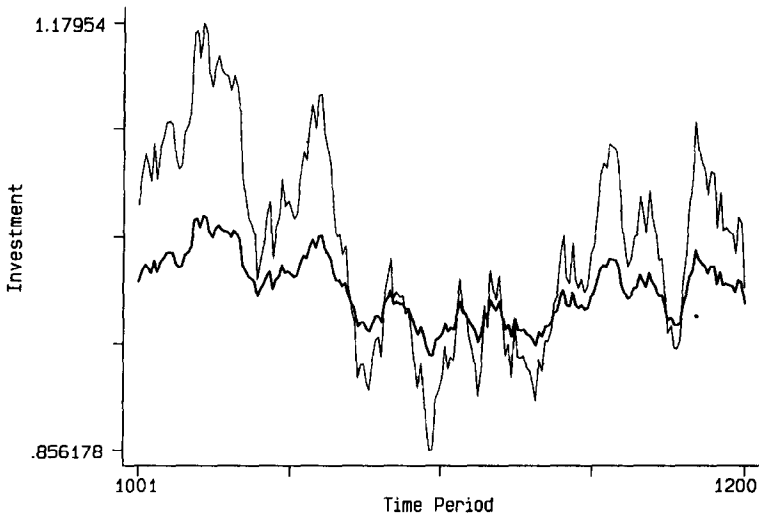
Fig. 3. Investment of sophisticated agents (thin) and unsophisticated agents (thick) in competitive equilibrium.

is based on the equilibrium in which half of the agents are sophisticated and the other half unsophisticated, graphs typical investment time series for the two types of agents.[9] In general (also for other values of θ), the time series for investment of sophisticated agents display much greater volatility in percentage terms than those for unsophisticated agents.[10] Sophisticated agents also choose higher volatility for income time series than do unsophisticated agents, but lower volatility for consumption time series than do unsophisticated agents.[11] The differences in volatilities, however, are the most pronounced for investment.

As a consequence of these differences in volatilities, aggregate time series in an economy in which all agents use the unsophisticated rule behave very differently from aggregate time series in the standard case where all agents use the sophisticated rule. Most notably, when all agents use the unsophisticated (constant-savings-rate) rule, time series for consumption, investment, and

---

[9] The thicker (darker) line in the graph is the investment time series for unsophisticated agents; the thinner (lighter) line is the investment time series for sophisticated agents.

[10] Specifically, we measure volatility as the standard deviation of the natural logarithm of the given time series.

[11] It is also the case that the volatilities of investment and output time series for sophisticated agents increase markedly as the equilibrium fraction of sophisticated agents falls. On the other hand, the volatility of consumption time series for sophisticated agents remains roughly constant as the equilibrium fraction of sophisticated agents falls.

income display equal volatilities (this follows from the fact that consumption, investment, and income are proportional to each other for the unsophisticated agent). Thus the small costs introduced into the standard framework lead to equilibrium outcomes that have very different dynamics. One of the successes of the real business cycle framework, namely the ability of intertemporal optimization – consumption smoothing – to account for the relative volatilities of the observed consumption, investment, and income series in large ('closed') economies, is now less of a given: if we believe that there are costs of implementation similar to those described here, then in fact this model loses its power to explain these aspects of the data.

Finally, Table 1 presents the sensitivity analysis with respect to v. It shows that more curvature in the utility function u leads to higher cutoff costs: with more curvature, the agent perceives volatile consumption as more painful, and hence the agent is less willing to use a rule of thumb which does not smooth consumption relative to income. However, the cutoff costs for the case $v = 5$ remain quite small: no consumers choose the unrestricted rule if its relative cost exceeds 0.079% of per period consumption (or \$16 per year assuming annual consumption expenditures of \$20,000). Table 1 also shows some properties of the aggregate time series for the alternative values of v and for different fractions of sophisticated agents in equilibrium; these are all qualitatively similar (and when all agents are rule-of-thumb decision makers, they are virtually identical).[12]

3.2. Other simple rules

In this section we replicate the results of the previous subsection using other rules of thumb in place of the constant-savings-rate rule. In particular, we look at three other rules of thumb:

(1) A partial-adjustment-of-capital rule: $k_{t+1} = \gamma k_t + (1 - \gamma)k_t^*$, where

$$k_t^* \equiv \left[\frac{\exp(z_t)\alpha\beta}{1 - \beta(1 - \delta)} \right]^{1/(1 - \alpha)}$$

and γ is chosen optimally (k_t^* is the steady state capital stock if the shock z_t remains indefinitely at its current level).

(2) A consumption function reading: $c_t = ak_t$, with a chosen optimally.

---

[12] For the case $v = 2$, Table 4 displays the decision rules of sophisticated agents and the savings rates of unsophisticated agents for a range of values of θ.

Table 1
Cutoff costs and aggregate time series statistics: Unrestricted rule vs. constant-savings-rate rule

$v = 1$

Cutoff costs

A	B	C	B/C
0.0064%	0.0162%	0.0094%	1.72

Aggregate time series statistics

θ	σ_c	σ_x	σ_y	cor(c, y)	cor(x, y)	cor(y, y_{-1})	cor(y, y_{-2})
1.0	2.43%	6.09%	3.08%	0.928	0.901	0.974	0.949
0.5	2.43%	4.47%	2.87%	0.982	0.955	0.970	0.941
0.0	2.61%	2.61%	2.61%	1.000	1.000	0.963	0.929

$v = 2$

Cutoff costs

A	B	C	B/C
0.013%	0.030%	0.018%	1.70

Aggregate time series statistics

θ	σ_c	σ_x	σ_y	cor(c, y)	cor(x, y)	cor(y, y_{-1})	cor(y, y_{-2})
1.0	2.35%	6.62%	3.25%	0.947	0.942	0.977	0.954
0.5	2.39%	4.80%	2.96%	0.990	0.978	0.972	0.945
0.0	2.61%	2.61%	2.61%	1.000	1.000	0.963	0.929

$v = 3$

Cutoff costs

A	B	C	B/C
0.021%	0.046%	0.028%	1.64

Aggregate time series statistics

θ	σ_c	σ_x	σ_y	cor(c, y)	cor(x, y)	cor(y, y_{-1})	cor(y, y_{-2})
1.0	2.32%	7.14%	3.40%	0.954	0.959	0.979	0.958
0.5	2.37%	5.10%	3.04%	0.993	0.987	0.973	0.948
0.0	2.61%	2.61%	2.61%	1.000	1.000	0.963	0.929

Table 1 continued

$v = 5$

Cutoff costs

A	B	C	B/C
0.039%	0.079%	0.051%	1.53

Aggregate time series statistics

θ	σ_c	σ_x	σ_y	cor(c, y)	cor(x, y)	cor(y, y_{-1})	cor(y, y_{-2})
1.0	2.30%	8.05%	3.66%	0.962	0.974	0.982	0.964
0.5	2.34%	5.60%	3.16%	0.997	0.995	0.975	0.952
0.0	2.61%	2.61%	2.61%	1.000	1.000	0.963	0.929

v is the coefficient of relative risk aversion ($v = 1$ denotes logarithmic preferences). The number in column A is η; below this cutoff cost, all agents use the unrestricted rule. The number in column B is $\bar{\eta}$; above this cutoff cost, all agents use the constant-savings-rate rule. The number in column C is the cutoff cost in a centralized, or decision-theoretic, setup; this is the cutoff cost above which all agents use the constant-savings-rate rule in an environment in which all agents are required to use the same rule. All cutoff costs are expressed in terms of a uniform percentage decrease in per period consumption. B/C is the ratio of column B to column C. θ is the competitive equilibrium fraction of sophisticated agents (i.e., agents using the unrestricted rule); $1 - \theta$ is the fraction of unsophisticated agents (i.e., agents using the constant-savings-rate rule). σ_i is the standard deviation of the logarithm of time series i (expressed as a percentage); cor(i, j) is the contemporaneous correlation between the logarithm of time series i and the logarithm of time series j; cor(y, y_{-k}) is the correlation between log output and log output lagged k periods ('c' denotes aggregate consumption consumption, 'x' denotes aggregate investment, and 'y' denotes aggregate output). Time series statistics are estimates based on simulated time series with 50,000 observations.

(3) A constant-capital rule: $k_t = \bar{k}$, where

$$\bar{k} \equiv \left[\frac{\alpha\beta}{1 - \beta(1 - \delta)} \right]^{1/(1-\alpha)}$$

is the deterministic steady state capital stock in the economy without decision rule costs.

Note that all these rules are very simple in that they each have at most one free parameter (the constant-capital rule has no free parameters). The partial-adjustment rule, like the constant-savings-rate rule, has roots in the traditional macroeconomic literature, but it offers much more in terms of possibilities for consumption smoothing. The constant-capital rule does the reverse: it smooths investment completely and makes consumption more volatile than income.

Finally, the $c = ak$ rule is similar to the constant-savings-rate rule, but offers much more in terms of consumption smoothing possibilities. Two of the four rules that we consider in this paper – the constant-savings-rate rule and the partial-adjustment rule – have built-in responses to the aggregate productivity shock; the others do not.

Besides having some intuitive appeal, we choose this set of rules to make our points as clearly as possible. First, the rules are all very simple. Second, they all appear in equilibrium except when the cost of the sophisticated rule is extremely low. Third, they give rise to quite different aggregate time series behavior. We do not consider more complicated rules, since this is unlikely to lead to any surprises: the greater the number of free parameters in the rules and the more complicated the rules are allowed to look, the 'better' will the rules be for the agents and the more easily will agents be willing to use the rules. A good example of this is the rule making consumption a log-linear function of the agent's own capital, aggregate capital, and the shock: this rule is known to be very close to the fully optimal rule, not only in utility terms but also in its time series properties (see, for example, Christiano, 1990).

It should also be pointed out that consumption rules that are more extreme than the constant-capital rule can be difficult to use. For example, a constant-consumption rule is not feasible unless consumption is set at a very low level, namely that level feasible with the worst possible sequence of realizations of the productivity shock – any higher constant consumption level would have to be violated with positive probability. Therefore, consumption has to be set in relation to at least one of the state variables in the economy.

Turning to the results, Table 2 shows a comparison of the three rules of thumb listed above (assuming $v = 2$ throughout).[13] We see that the cutoff costs above which all agents use the rule of thumb in equilibrium are all very low, with the highest value recorded for the constant-capital rule: 0.089% of per period consumption. Furthermore, it is clear that partial adjustment of capital provides for substantial consumption smoothing; a similar proposition holds for the $c = ak$ rule, but to a somewhat lesser extent. The 'best' rule – in the sense that this rule has the smallest cutoff cost above which all agents use the rule in equilibrium – is the partial-adjustment rule. The aggregate time series statistics when all agents use either the partial-adjustment or $c = ak$ rules are much closer to those of the standard model than when all agents use either the constant-savings-rate or constant-capital rules.

Finally, we also let two rules of thumb compete, assuming implicitly that the fully sophisticated rule is prohibitively costly for any agent. These experiments

---

[13] The optimal choice for γ in the partial-adjustment rule varies only to a small degree with the value of η: the optimal choice is close to 0.955 in all cases. Similarly, the optimal choice for a in the $c = ak$ rule is close to 0.0725 regardless of the value of η.

Table 2
Cutoff costs and aggregate time series statistics: Unrestricted rule vs. three rules of thumb

Partial-adjustment-of-capital rule

Cutoff costs

A	B	C	B/C
0.0019%	0.0142%	0.0043%	3.26

Aggregate time series statistics

θ	σ_c	σ_x	σ_y	cor(c, y)	cor(x, y)	cor(y, y_{-1})	cor(y, y_{-2})
1.0	2.35%	6.62%	3.25%	0.947	0.942	0.977	0.954
0.5	2.39%	5.95%	3.10%	0.948	0.927	0.974	0.949
0.0	2.49%	5.18%	2.90%	0.939	0.878	0.971	0.942

c = ak rule

Cutoff costs

A	B	C	B/C
0.0087%	0.0346%	0.0143%	2.42

Aggregate time series statistics

θ	σ_c	σ_x	σ_y	cor(c, y)	cor(x, y)	cor(y, y_{-1})	cor(y, y_{-2})
1.0	2.35%	6.62%	3.25%	0.947	0.942	0.977	0.954
0.5	2.44%	6.53%	3.15%	0.913	0.896	0.975	0.951
0.0	2.63%	6.49%	3.01%	0.860	0.798	0.973	0.945

Constant-capital rule

Cutoff costs

A	B	C	B/C
0.031%	0.089%	0.046%	1.92

Aggregate time series statistics

θ	σ_c	σ_x	σ_y	cor(c, y)	cor(x, y)	cor(y, y_{-1})	cor(y, y_{-2})
1.0	2.35%	6.62%	3.25%	0.947	0.942	0.977	0.954
0.5	2.49%	3.66%	2.77%	0.997	0.989	0.968	0.937
0.0	3.02%	0.00%	2.24%	1.000	0.000	0.950	0.903

The coefficient of relative risk aversion $v = 2$ for all three cases. θ is the competitive equilibrium fraction of agents using the unrestricted rule; $1 - \theta$ is the fraction of agents using the rule of thumb. For additional explanation, see the notes to Table 1.

are designed to show more explicitly that the aggregate time series behavior may depend crucially not only on the costs of the rules of thumb relative to those of unrestricted behavior, but also on the relative costs of different simple rules. Table 3 makes this point clear, contrasting the constant-savings-rate rule (implying no consumption smoothing) with the $c = ak$ rule (implying substantial consumption smoothing). Table 3 also contrasts the partial-adjustment rule with, respectively, the constant-savings-rate rule and the $c = ak$ rule.

3.3. Socially optimal allocations

The environment studied here has externalities in that the consumption possibility set depends on the aggregate law of motion of the distribution of capital. In particular, when the agents make decisions, they do not internalize the effect of their choices on the consumption possibilities for agents who follow rules of thumb: *aggregate* investment at time t decisively affects how these agents will consume at time $t + 1$.

Given that all agents are alike as of time 0, the planning problem is straightforward: maximize total utility of all agents. This, in turn, is equivalent to maximizing the weighted sum of sophisticated and unsophisticated agents' utilities, with the weights being the fraction of each type. That fraction is now a direct choice variable from the point of view of the planner. The constraints to the planning problem are also clear; they are simply the aggregate technology constraints and the restriction that the two kinds of agents have to choose consumption processes in $\mathscr{C}$ and $\mathscr{C}_u$, respectively.

To gain further insight into the externality described above, we solve the planning problem numerically for the economy in which agents choose between the unrestricted rule and the constant-savings-rate rule.[14] We find that the range of values of the effort cost η for which both types of agents coexist is the same as for the competitive equilibrium allocation. Furthermore, as shown in Smith (1992), over the range where there is only one type of agent, the socially optimal decision rules coincide with the competitive equilibrium decision rules – a first welfare theorem for special versions of this environment.[15] In other words, in the case where all agents are of the same type, the agents choose socially optimal decision rules in equilibrium, and the externality disappears.

However, over the range where θ, the competitive equilibrium fraction of sophisticated agents, is in (0, 1), the competitive equilibrium produces suboptimal behavior. First, decision rules and savings rates are not optimal, even if the

---

[14] Once again, the qualitative results reported here carry over to the economies in which agents choose between other types of rules.

[15] It is straightforward to verify that the first-order conditions in the two cases coincide and that these conditions are sufficient for maxima in both cases.

Table 3
Cutoff costs and aggregate time series statistics: Three pairs of rules of thumb

Partial-adjustment-of-capital rule vs. c = ak rule

Cutoff costs

A	B	C	B/C
0.0093%	0.0107%	0.0100%	1.07

Aggregate time series statistics

θ	σ_c	σ_x	σ_y	cor(c, y)	cor(x, y)	cor(y, y_{-1})	cor(y, y_{-2})
1.0	2.49%	5.18%	2.90%	0.939	0.878	0.971	0.942
0.5	2.54%	5.81%	2.95%	0.903	0.837	0.972	0.944
0.0	2.63%	6.49%	3.01%	0.860	0.798	0.973	0.945

Partial-adjustment-of-capital rule vs. Constant-savings-rate rule

Cutoff costs

A	B	C	B/C
0.012%	0.015%	0.014%	1.12

Aggregate time series statistics

θ	σ_c	σ_x	σ_y	cor(c, y)	cor(x, y)	cor(y, y_{-1})	cor(y, y_{-2})
1.0	2.49%	5.18%	2.90%	0.939	0.878	0.971	0.942
0.5	2.50%	3.85%	2.76%	0.984	0.943	0.967	0.936
0.0	2.61%	2.61%	2.61%	1.000	1.000	0.963	0.929

c = ak rule vs. Constant-savings-rate rule

Cutoff costs

A	B	C	B/C
0.0021%	0.0052%	0.0036%	1.43

Aggregate time series statistics

θ	σ_c	σ_x	σ_y	cor(c, y)	cor(x, y)	cor(y, y_{-1})	cor(y, y_{-2})
1.0	2.63%	6.49%	3.01%	0.860	0.798	0.973	0.945
0.5	2.51%	4.37%	2.80%	0.964	0.898	0.968	0.938
0.0	2.61%	2.61%	2.61%	1.000	1.000	0.963	0.929

The coefficient of relative risk aversion $v = 2$ for all three cases. θ is the competitive equilibrium fraction of agents using the the first rule of thumb; $1 - \theta$ is the fraction of agents using the second rule of thumb. For additional explanation, see the notes to Table 1.
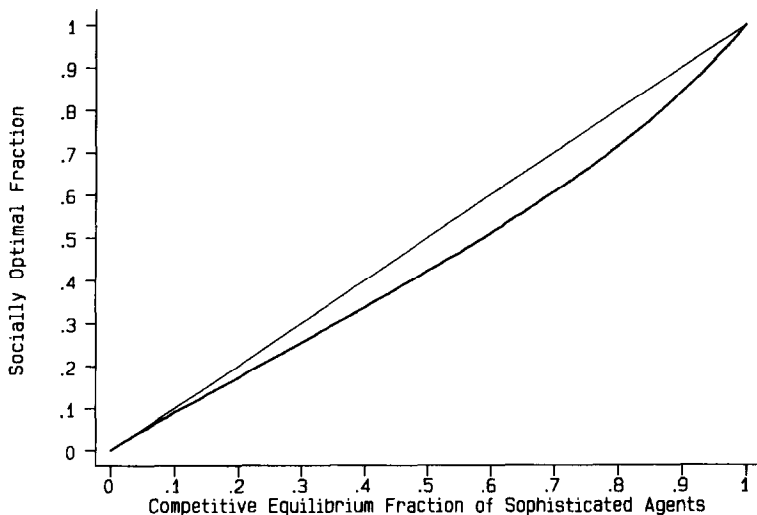
Fig. 4. Socially optimal fraction vs. competitive equilibrium fraction of sophisticated agents.

agents are not able to select among rules of differing sophistication. Numerically, the differences between the competitive equilibrium and socially optimal savings rates are small. However, the socially optimal decision rules for sophisticated agents differ noticeably from the decision rules that sophisticated agents choose in competitive equilibrium.

Second, the competitive equilibrium fraction of sophisticated agents, given a value of the effort cost η, is not in general optimal. We find that the optimal allocation dictates a (weakly) smaller number of sophisticated agents than does the competitive equilibrium. This result is robust to changes in the values of parameters, including the cost η. Fig. 4 graphs the socially optimal fraction of sophisticated agents against that of the corresponding competitive equilibrium for the case of the constant-savings rate rule.

Table 4 displays the savings rates of unsophisticated agents and the decision rules of sophisticated agents for a range of values of θ, the competitive equilibrium fraction of sophisticated agents. For each of these values of θ, Table 4 also displays the savings rate and the decision rule (as well as the socially optimal fraction θ^* of sophisticated agents) that the social planner chooses.

We also note that when the competitive equilibrium fails to be optimal, i.e., for $\eta \in (\underline{\eta}, \bar{\eta})$, the social planner would also be able to increase utility (*net* of effort cost) for all agents. Thus, moving from the competitive equilibrium to the socially optimal allocations is also (ex post) Pareto improving. This can be illustrated by the following facts. Let the utility (gross of effort cost) of an agent of type $i \in \{S, U\}$ in the competitive and socially optimal allocations for a given

Table 4
Competitive equilibrium and socially optimally decision rules: Unrestricted rule vs. constant-savings-rate rule

	θ	Unsophisticated agents s	Sophisticated agents D_0	D_1	D_2	D_3
Competitive eq.	1.000	0.25698	0.0854	0.9765	0.0000	0.0720
Socially optimal	1.000	—	0.0854	0.9765	0.0000	0.0720
Competitive eq.	0.900	0.25679	0.0854	0.9767	− 0.0002	0.0720
Socially optimal	0.842	0.25693	0.0863	0.9787	− 0.0024	0.0732
Competitive eq.	0.750	0.25668	0.0859	0.9774	− 0.0010	0.0724
Socially optimal	0.658	0.25687	0.0882	0.9816	− 0.0059	0.0749
Competitive eq.	0.500	0.25661	0.0876	0.9802	− 0.0043	0.0740
Socially optimal	0.422	0.25678	0.0904	0.9858	− 0.0106	0.0777
Competitive eq.	0.250	0.25658	0.0918	0.9857	− 0.0110	0.0774
Socially optimal	0.214	0.25667	0.0936	0.9906	− 0.0164	0.0810
Competitive eq.	0.100	0.25656	0.0978	0.9917	− 0.0186	0.0816
Socially optimal	0.090	0.25659	0.1012	0.9951	− 0.0229	0.0841
Competitive eq.	0.000	0.25651	0.1105	0.9999	− 0.0303	0.0880
Socially optimal	0.000	0.25651	—	—	—	—

The coefficient of relative risk aversion $v = 2$. θ is the fraction of sophisticated agents (i.e., the fraction using the unrestricted rule); $1 - \theta$ is the fraction of unsophisticated agents (i.e., the fraction using the constant-savings-rate rule). Rows labelled 'Competitive eq.' contain the decision rule choices of both types of agents in competitive equilibrium. Rows labelled 'Socially optimal' contain the socially optimal fraction of sophisticated agents and the socially optimal decision rule choices for sophisticated and unsophisticated agents, given the corresponding competitive equilibrium fraction. For unsophisticated agents, s is the optimal savings rate. The decision rule used by sophisticated agents is given by $\log(K_{St+1}) = D_0 + D_1 \log(K_{St}) + D_2 \log(K_{Ut}) + D_3 z_t$, where K_{St} and K_{Ut} are the capital stocks held by, respectively, sophisticated and unsophisticated agents and z_t is the aggregate productivity shock.

η be, respectively, $U_i(\theta)$ and $U_i^*(\theta^*)$, with the asterisk (*) denoting the optimal allocation. Numerical calculations reveal that $U_i(\theta) < U_i^*(\theta^*)$ for both values of i. Now, since $\theta^* < \theta$, the planner wishes to reallocate the fraction $\theta - \theta^*$ of agents from sophisticated to unsophisticated decision rules. These agents are better off since $U_U^*(\theta^*) > U_S(\theta) - \eta = U_U(\theta)$, where the inequality is the same as the one above and the equality is a condition that holds in competitive equilibrium. Next, the fraction $1 - \theta$ of agents continue to use unsophisticated decision rules; these agents are better off because $U_U^*(\theta^*) > U_U(\theta)$. Similarly, the fraction θ^* of agents continue to use sophisticated decision rules; these agents are also better off in the new allocation: $U_S^*(\theta^*) - \eta > U_S(\theta) - \eta$.

Finally we note that the increases in utility that are possible in going from the competitive to the socially optimal allocation are numerically very small. The aggregate (average) increase in welfare, which varies with the cost η, is largest when the cost η is such that the competitive equilibrium fraction of sophisticated agents is near one-half. Even in this case, however, the aggregate welfare gain, measured in terms of a uniform percentage increase in per period consumption, is only 0.0001% when v (the coefficient of relative risk aversion) equals one and 0.0004% when v equals five.

4. Conclusions

How should our results be interpreted? Clearly, two quite different interpretations are available. First, one might say that since the constant-savings-rate rule (and other very simple rules of thumb as well) make the aggregate behavior of the standard model look worse – aggregate consumption smoothing is weakened – we should conclude that, in the real world, the costs of sophisticated behavior must be very low (or, alternatively, that the constant-savings-rate and constant-capital rules must be expensive relative to other simple rules which do lead to aggregate consumption smoothing). In other words, since the aggregate behavior of the model is quite sensitive to decision rule costs, the relative magnitudes of these costs can be 'estimated' with high accuracy, in particular to be smaller than the relevant cutoffs calculated in the text.

Second, at least to the extent that one finds acceptable the way in which this paper models unsophistication, the results can be interpreted as alarming for all those who use the stochastic, representative-agent growth model as a basis for macroeconomic theorizing. This view is based on the belief that macroeconomic models should have quantitatively reasonable microeconomic underpinnings or, put differently, that a good model should fit not only the macro but also the micro data. According to this view, the real-world costs facing consumers are in fact higher than the cutoffs we report in this paper, so that reasonable microeconomic underpinnings should include these costs.

We are inclined toward the second interpretation, but we feel that several caveats are necessary. First, there is clearly a need to solve the conceptual problems of finding a way to formalize costly decision making. Our approach here is ad hoc, and it remains to be seen whether our results would hold up in structures with a deeper, and less restrictive, modeling structure. A more structural approach seems especially important in light of one of the key findings of this paper: the relative costs of different rules of thumb, about which we have very poor information, seem crucial in order to determine the model's aggregate behavior.

Second, we also think that further work on analyzing the sensitivity of the intertemporal optimization framework to various costs of behavioral

sophistication needs to tackle a couple of other issues which we have not been able to deal with in this paper. One is that our structure is not recursive; one would, ideally, recognize that agents face the tradeoff between sophisticated and unsophisticated behavior at every point in time, and not just at time zero. A related problem is that of agent heterogeneity: the economic model in this paper is as close as one could get to the representative-agent framework. If agents are not the same ex ante, and if they face uninsurable idiosyncratic risk in addition to aggregate productivity shocks, then the quantitative analysis may provide a different answer.[16] The reason for this, of course, is that the idiosyncratic risk is much larger than the aggregate risk. With higher variance of the risk, the costs of not being sophisticated go up, and this may change the results significantly. However, the difficulty of analyzing models with agent heterogeneity and aggregate risk is well known (see, for example, Díaz-Giménez, Prescott, Fitzgerald, and Alvarez, 1992), and economists have only just begun the search for solution techniques to be applied to this class of models.

Appendix

This appendix describes the numerical algorithm that is used to compute the competitive equilibrium defined by Definition 1. The algorithm is described in the context of the economy in which agents choose between an unrestricted rule and the constant-savings-rate rule. It is straightforward to modify the algorithm to handle other pairs of rules.

The numerical algorithm makes use of linear-quadratic approximations in order to reduce computational burden. The key idea of the numerical algorithm is to fix a fraction θ of sophisticated agents, solve for the competitive equilibrium decision rule choices of sophisticated and unsophisticated agents given θ, and then calculate the cost $\eta = U_S(\theta) - U_U(\theta)$ that supports θ as a competitive equilibrium.

The first step is to solve the sophisticated agents' problem (PS) and the unsophisticated agents' problem (PU), given a fraction θ of sophisticated agents. Note that the equilibrium pricing functions $r_t = MPK_t$ and $w_t = MPL_t$ can be inserted into the constraints of problems PS and PU without loss of generality, since individual agents take as given the behavior of aggregate variables.

First consider the optimization problem faced by a sophisticated agent, given θ. Recalling that $K_t = \theta K_{St} + (1 - \theta)K_{Ut}$, the constraints to problem PS can be used to express c_t as a function of the period t state variables $\tilde{k}_t \equiv \log(k_t)$, $\tilde{K}_{St} \equiv \log(K_{St})$, $\tilde{K}_{Ut} \equiv \log(K_{Ut})$, and z_t and the period t choice variable $\tilde{k}_{t+1}$. Using this expression for c_t, the period t utility function $u(c_t)$ can be

---

approximated by a quadratic function in a neighborhood of the steady state values of the period t state and choice variables: $u(c_t) \approx \hat{u}(\tilde{k}_t, \tilde{K}_{St}, \tilde{K}_{Ut}, \bar{z}_t, \tilde{k}_{t+1}) \equiv \xi_t' Q_1 \xi_t + 2\tilde{K}_{St} Q_2 \xi_t + 2\tilde{k}_{t+1} W_1 \xi_t + R_1 \tilde{k}_{t+1}^2 + 2W_2 \tilde{k}_{t+1} \tilde{K}_{St} + S\tilde{K}_{St}^2$, where $\xi_t = [1 \; \tilde{k}_t \; \tilde{K}_{Ut} \; z_t]'$.

In order that the agent's problem be well-defined, it is necessary to assign to the sophisticated agent a perceived law of motion for the aggregate capital stock $\tilde{K}_{Ut}$ held by unsophisticated agents. The decision rule used by a typical unsophisticated agent is given by $k_{t+1} = f_U(k_t, K_{St}, K_{Ut}, z_t)$, where the specific form of the function f_U is given in Section 2.2. This decision rule can be reexpressed in the form $\tilde{k}_{t+1} = \tilde{f}_U(\tilde{k}_t, \tilde{K}_{St}, \tilde{K}_{Ut}, z_t)$, where $\tilde{f}_U(a, b, c, \cdot) \equiv \log f_U(\exp(a), \exp(b), \exp(c), \cdot)$. The function $\tilde{f}_U$ can in turn be approximated by a linear function of its arguments in a neighborhood of the steady state values of $\tilde{k}_t$, $\tilde{K}_{St}$, $\tilde{K}_{Ut}$, and z_t: $\tilde{k}_{t+1} = e_0 + e_1 \tilde{k}_t + e_2 z_t + e_3 \tilde{K}_{St} + e_4 \tilde{K}_{Ut}$. It turns out that $e_0 = \log D - (C/D)\tilde{K}$, $e_1 = C/D$, $e_2 = (s/D)\exp(\alpha \bar{K})$, and $e_3 = e_4 = 0$, where s is the savings rate used by the unsophisticated agent, $\bar{K}$ is the (common) steady state value of $\tilde{k}_t$, $\tilde{K}_{St}$, and $\tilde{K}_{Ut}$, $C = s\alpha \exp(\alpha \bar{K}) + (1 - \delta)\exp(\bar{K})$, and $D = s \exp(\alpha \bar{K}) + (1 - \delta)\exp(\bar{K})$. Note that the coefficients e_0, e_1, and e_2 are all functions of the savings rate used by the unsophisticated agent. Since, in equilibrium, all unsophisticated agents use the same savings rate, it follows that $K_{Ut+1} = E_0 + E_1 K_{Ut} + E_2 z_t$, where $E_i = e_i$, $i = 0, 1, 2$. (When solving the unsophisticated agent's problem PU below, it is useful to maintain a distinction between the coefficients e_i used by a given small unsophisticated agent and the coefficients E_i used by all other unsophisticated agents.) Hence, the coefficients E_i are functions of the (common) savings rate s used by unsophisticated agents. The sophisticated agent takes s (and hence the coefficients E_i) as given when solving problem PS.

Given these perceived laws of motion, the linear-quadratic approximation to problem PS can be written in recursive form as follows:

$$\xi_t' P \xi_t = \max_{\tilde{k}_{t+1}} \left[\hat{u}(\tilde{k}_t, \tilde{K}_{St}, \tilde{K}_{Ut}, z_t, \tilde{k}_{t+1}) + \beta \xi_{t+1}' P \xi_{t+1} \right]$$

subject to

$$\xi_{t+1} = A_1 \xi_t + B_1 \tilde{k}_{t+1},$$

where

$$A_1 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ E_0 & 0 & E_1 & E_2 \\ 0 & 0 & 0 & \rho \end{bmatrix},$$

$$B_1 = [0 \; 1 \; 0 \; 0]',$$

and P is an unknown matrix that must be computed. When solving this problem, the agent takes $\tilde{K}_{St}$ as given. Uncertainty can be suppressed in this problem without loss of generality by the certainty equivalence principle.

Substituting the constraint into the objective function above, the first-order condition (given P) is

$$\tilde{k}_{t+1} = F_1 \xi_t + F_2 \tilde{K}_{St}, \tag{1}$$

where

$$F_1 = -(R_1 + \beta B_1' P B_1)^{-1}(W_1 + \beta B_1' P A_1),$$
$$F_2 = -(R_1 + \beta B_1' P B_1)^{-1} W_2.$$

Eq. (1) is the decision rule of an individual sophisticated agent. Since, in equilibrium, all sophisticated agents choose the same decision rule, we can impose on the decision rule (1) the consistency condition $\tilde{k}_t = \tilde{K}_{St} = \Theta \xi_t$ (where $\Theta = [0\ 1\ 0\ 0]$) to yield $\tilde{k}_{t+1} = D \xi_t$, where $D = F_1 + \Theta F_2$. Given this solution for $\tilde{k}_{t+1}$ and the consistency condition $\tilde{K}_{St} = \Theta \xi_t$, the objective function above can be expressed as a quadratic form $\xi_t' \bar{P} \xi_t$ in the period t state vector ξ_t. We seek a matrix P such that $\bar{P} = P$. This matrix can be computed via an iterative process by setting $P^{(0)} = 0$, computing $\bar{P}$, setting $P^{(1)} = \bar{P}$, etc., until convergence is achieved.

The solution to the sophisticated agent's problem PS can be summarized by a function $H_S : E \rightarrow D$ that maps perceptions of unsophisticated agent's behavior (as determined by the decision rule coefficients E_i) into the optimal decision rule coefficients $D = [D_0\ D_1\ D_2\ D_3]$ chosen by sophisticated agents.

Next, consider the optimization problem faced by an unsophisticated agent, given θ. Write the quadratic approximation $\hat{u}$ to the period t utility function as follows:

$$u(c_t) \approx \hat{u}(\tilde{k}_t, \tilde{K}_{St}, \tilde{K}_{Ut}, z_t, \tilde{k}_{t+1}) = \zeta_t' Q \zeta_t + R_2 \tilde{k}_{t+1}^2 + 2\tilde{k}_t W \zeta_t,$$

where

$$\zeta_t = [1\ \tilde{k}_t, \tilde{K}_{St}, \tilde{K}_{Ut}, z_t]'.$$

The unsophisticated agent seeks to maximize (by choice of a state-contingent sequence $\{k_t\}_{t=1}^{\infty}$):

$$E_0 \sum_{t=0}^{\infty} \beta^t \hat{u}(k_t, K_{St}, K_{Ut}, z_t, k_{t+1}), \qquad \text{given} \quad \zeta_0,$$

subject to

$$\zeta_{t+1} = A_2 \zeta_t + B_2 k_t + \varepsilon_{t+1},$$

where

$$E(v_{t+1}) = 0, \qquad E(v_{t+1} v'_{t+1}) = \Sigma,$$

$$A_2 = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ D_0 & 0 & D_1 & D_2 & D_3 \\ E_0 & 0 & 0 & E_1 & E_2 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix},$$

$$B_2 = [0 \ 1 \ 0 \ 0 \ 0]'.$$

Σ is a 5×5 matrix of zeros save for σ_ε^2 as the last diagonal element. The coefficients D_i and E_i in the matrix A_2 reflect the unsophisticated agent's perceived laws of motion for the aggregate state variables $\tilde{K}_{St}$ and $\tilde{K}_{Ut}$.

Unlike the sophisticated agent, the unsophisticated agent cannot freely choose a sequence $\{k_t\}$ to maximize his/her objective function, but is constrained to use a constant savings rate rule, with savings rate s chosen optimally. As discussed above, a first-order approximation to this decision rule can be written $\tilde{k}_t = F\zeta_t$, where $F = [e_0 \ e_1 \ 0 \ 0 \ e_2]$ and the coefficients e_i can be expressed as functions of the savings rate s chosen by the small unsophisticated agent. Given F (as determined by s) and an initial condition ζ_0, the value of the unsophisticated agent's objective function is given by

$$\text{trace}\left[(Q + F'R_2F + 2F'W) \sum_{t=0}^{\infty} \beta^t \, E_0(\zeta_t, \zeta'_t) \right], \qquad (2)$$

where the recursive relation

$$E_0(\zeta_{t+1} \zeta'_{t+1}) = (A_2 + B_2F) \, E_0(\zeta_t \zeta'_t) (A_2 + B_2F)' + \Sigma \qquad (3)$$

can be used to evaluate the conditional expectation in Eq. (2). Eqs. (2) and (3) can be used in a conventional hillclimbing algorithm to determine the optimal savings rate.

The solution to the unsophisticated agent's problem PU can be summarized by a function $H_U: (D, E) \to e$ that maps perceptions of the behavior of sophisticated agents and of *other* unsophisticated agents (as determined by the decision rule coefficients D_i and E_i, respectively) into the optimal decision rule coefficients e_i chosen by an unsophisticated agent. In equilibrium, unsophisticated agents are identical, so that $e_i = E_i$.

A competitive equilibrium in the linear-quadratic economy is a pair of decision rule coefficients (D^*, E^*) such that $D^* = H_S(E^*)$ and $E^* = H_U(D^*, E^*)$.

The fixed point (D^*, E^*) can be computed iteratively beginning from some initial pair of coefficients (D^0, E^0).

The solution to the social planner's problem can be computed in two steps. In the first step, the fraction of sophisticated agents θ is fixed and linear-quadratic methods analogous to those used above to solve the unsophisticated agent's problem are used to compute the optimal decision rule for sophisticated agents and the optimal savings rate for unsophisticated agents. In particular, in the linear quadratic approximation to the social planner's problem (given θ), the social planner chooses a pair of decision rules $\tilde{K}_{St+1} = F_S \omega_t$ and $\tilde{K}_{Ut+1} = F_U \omega_t$, where $\omega_t = [1 \ \tilde{K}_{St} \ \tilde{K}_{Ut} \ z_t]'$. The social planner can freely choose the elements of F_S. The elements of F_U, on the other hand, are functions of the savings rate s chosen by the social planner for use by unsophisticated agents (see the discussion above). Given θ, the social planner chooses F_S and s so as to maximize the quadratic approximation to the appropriate objective function, i.e., the weighted average of sophisticated agents' and unsophisticated agents' utilities, with the weights given by the fraction of each type of agent in the economy. In the second step, the social planner searches over values of θ so as to maximize the quadratic approximation to the planner's objective function, taking into account the mapping from θ to the corresponding socially optimal choices of F_S and s.

References

Akerlof, G.A. and J.L. Yellen, 1985a, Can small deviations from rationality make significant differences to economic equilibria?, American Economic Review 75, 823–838.

Akerlof, G.A. and J.L. Yellen, 1985b, A near-rational model of the business cycle, with wage and price inertia, Quarterly Journal of Economics 100, Suppl., 823–838.

Aumann, R., 1992, Perspectives on bounded rationality, in: Y. Moses, ed., Theoretical aspects of reasoning about knowledge: Proceedings of the fourth conference (TARK 1992), Manuscript (Weizmann Institute, Jerusalem).

Brock, W.A. and L.J. Mirman 1972, Optimal economic growth and uncertainty: The discounted case, Journal of Economic Theory 4, 479–513.

Caballero, R.J., 1991, Fuzzy preferences and aggregate consumption, Discussion paper no. 520 (Columbia University, New York, NY).

Campbell, J.Y. and N.G. Mankiw, 1990, Permanent income, current income, and consumption, Journal of Business and Economic Statistics 8, 265–279.

Christiano, L.J., 1990, Linear-quadratic approximation and value-function iteration, Journal of Business and Economic Statistics 8, 99–113.

Cochrane, J.H., 1989, The sensitivity of tests of the intertemporal allocation of consumption to near-rational alternatives, American Economic Review 79, 319–337.

Crettez, B. and P. Michel, 1990, Economically rational expectations equilibrium, Discussion paper no. 9055 (CORE, Louvain-la-Neuve).

Díaz-Giménez, J., E.C. Prescott, T. Fitzgerald, and F. Alvarez, 1992, Banking in computable general equilibrium economies, Journal of Economic Dynamics and Control 16, 533–559.

De Long, J.B., A. Shleifer, L.H. Summers, and R.J. Waldmann, 1990, Noise trader risk in financial markets, Journal of Political Economy 98, 703–738.

Evans, G.W. and S. Honkapohja, 1993, Adaptive learning and expectational stability: An introduction, in: A. Kirman and M. Salmon, eds., Learning and rationality in economics (Basil Blackwell, Oxford).

Evans, G.W. and G. Ramey, 1992, Expectation calculation and macroeconomic dynamics, American Economic Review 82, 207–224.

Haltiwanger, J. and M. Waldman, 1985, Rational expectations and the limits of rationality: An analysis of heterogeneity, American Economic Review 75, 326–340.

Haltiwanger, J. and M. Waldman, 1991, Responders versus non-responders: A new perspective on heterogeneity, Economic Journal 101, 1085–1102.

Hansen, G.D. and R. Wright, 1992, The labor market in real business cycle theory, Federal Reserve Bank of Minneapolis Quarterly Review, 2–12.

Hurwicz, L., S. Reiter, and D. Saari, 1994, Computational complexity and communication in resource allocation mechanisms (Cambridge University Press, Cambridge, under contract).

İmrohoroğlu, S., 1989, The cost of business cycles with indivisibilities and liquidity constraints, Journal of Political Economy 97, 1364–1383.

Ingram, B.F., 1990, Equilibrium modeling of asset prices: Rationality versus rules of thumb, Journal of Business and Economic Statistics 8, 115–126.

Lipman, B.L., 1991, How to decide how to decide how to … : Modeling limited rationality, Econometrica 59, 1105–1126.

Lucas, R.E. Jr., 1987, Models of business cycles (Basil Blackwell, New York, NY).

Mirrlees, J.A. and N.H. Stern, 1972, Fairly good plans, Journal of Economic Theory 4, 268–288.

Rosenthal, R.W., 1991, Rules of thumb in games, Manuscript (Boston University, Boston, MA).

Sargent, T.J., 1992, Bounded rationality in macroeconomics, Manuscript (Hoover Institution, Stanford, CA; University of Chicago, Chicago, IL).

Smith, A.A., Jr., 1990, Three essays on the solution and estimation of dynamic macroeconomic models, Unpublished Ph.D. thesis (Duke University, Durham, NC).

Smith, A.A., Jr., 1992, Near-rational behavior and the real business cycle, Manuscript (Carnegie Mellon University, Pittsburgh, PA).

Solow, Robert M., 1956, A contribution to the theory of economic growth, Quarterly Journal of Economics 70, 65–94.

Stokey, N.L. and R.E. Lucas, Jr., with E.C. Prescott, 1989, Recursive methods in economic dynamics (Harvard University Pres, Cambridge, MA).

Weil, P., 1991, Hand-to-mouth consumers and asset prices, Discussion paper no. 1562 (Harvard Institute of Economic Research, Cambridge, MA).