

Time Orientation and Asset Prices

Per Krusell, Burhanettin Kuruşçu, and
Anthony A. Smith, Jr.*

August 2001

Abstract

We analyze a general-equilibrium asset pricing model where a small subset of the consumers/investors have a short-run “urge to save”. That is, their attitude toward consumption in the long run is a standard one—they do place zero weight on consumption far enough out in the future—but their short-run effective rates of discount may be negative. Our model, which is an elaboration on the framework proposed by Faruk Gul and Wolfgang Pesendorfer, does not feature time inconsistencies. Thus, we view consumers as fully rational, but subject to specific “internal frictions” in the form of temptation and self-control problems. The model nests the Mehra-Prescott model and we use it as a way of interpreting the wealth and asset pricing data. Some aspects of these data, we argue, can possibly be better understood using our model than the standard one.

*Per Krusell and Burhanettin Kuruşçu are at the Department of Economics, Harkness Hall, University of Rochester, Rochester, NY 14627, telephone (716) 273-4903, fax (716) 271-3900, and email pekr@troi.cc.rochester.edu and kubu@troi.cc.rochester.edu, respectively; Krusell is also at the Institute for International Economic Studies and the CEPR; Anthony A. Smith, Jr. is at Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh, PA 15213, telephone (412) 268-7583, fax (412) 268-7357, email smithaa@andrew.cmu.edu. We would like to thank our discussant Qiang Dai as well as Larry Epstein for helpful comments and suggestions. Krusell and Smith thank the National Science Foundation for financial support.

1 Introduction

Recent psychological, sociological, and experimental evidence suggest that “time orientation”—a term for how people compare the present to the future—is a very important determinant of the intertemporal decisions that people make. In particular, time orientation is argued to differ a lot among consumers and even to embody “inconsistencies” as evidenced by so-called preference reversals: consumers seem to change their minds as time passes. The idea that many consumers, or even the average consumer, would have such a lack of “self-control” has recently gained support also in economics.

Self-control problems that take the form of “short-run urges” have recently received particular attention in the savings literature. In this paper, we consider such problems, but we focus on the opposite of what has typically been studied: rather than emphasizing a present-bias—an urge for immediate consumption—we hypothesize that an agent may have an urge to save. We present no independent evidence for this hypothesis, but we think it is not an implausible one for a subset of the population. A focus on wealth accumulation alongside frugal attitudes toward consumption is if not a typical so at least a not uncommon description of the behavior of many households. On the one hand, one could simply think of such behavior as originating from a standard utility function with a low discount rate. On the other hand, it is conceivable too that such consumers do see themselves consuming their wealth, but only “just not yet”, and that they do not change these attitudes over time. For example, a hard-to-explain feature of the savings data is the reluctance of many old-age consumers to dissave in cases where there are no natural heirs. That is, the interpretation would be that continued wealth accumulation—perhaps a focus on “not running down one’s assets”—also can be regarded as an urge, or a temptation.

On a general level, we consider preference heterogeneity among consumers as rather natural; in fact, although most of the macroeconomic literature works with a representative agent, we view this practice as more of a modelling short-cut than motivated by realism. In this paper, we additionally ask the reader to consider the possibility that *at least some* consumers have “urges”, and that these urges may take different expressions in different consumers. We also think such a supposition is quite natural, despite its possibly radical tone (at least in the context of the macroeconomic literature). Our basic argument on the level of motivation, then, continues by pointing out that it would tend to follow from these assumptions that those consumers

with a future-bias—an urge to save—will accumulate more wealth than the average consumer and thus become dominant asset holders in the economy. That is, those consumers who have a future-bias would turn out to be of particular interest for studies of the aggregate economy, both in terms of aggregate savings (and other macroeconomic variables) and in terms of the influence on asset prices that these agents would have.

The purpose of this paper is to formulate a simple equilibrium model with these features: we model an urge to save, and we embed consumers with such an urge in a general-equilibrium asset-pricing model. More precisely, we look at a model where a small set of consumers have a future-bias and where the remainder of the population have no future-bias (are standard, unbiased consumers or have a present-bias). The model allows an alternative interpretation of the puzzles in the asset pricing and wealth data. For example, we illustrate how a low risk-free rate is not surprising in our economy, and we explore the conditions on preferences under which both the risk-free rate and the equity premium take on realistic values. These illustrations are carried out in a model parameterized so that most of the wealth is held by very few investors—as in the data—and where the aggregate consumption movements nevertheless are significantly influenced also by the middle- and low-wealth households—as in the data.

Although the setup we use is in the spirit of the incomplete-markets models in Krusell and Smith (1998) and Guvenen (2001), we simplify quite a bit with respect to the former by assuming two types of agents, and with respect to the latter by assuming complete markets; we maintain the presence of a borrowing constraint.¹ We show that under certain assumptions on parameters, the borrowing constraint will bind for the poorer agents on the entire equilibrium path. This implies that this agent’s consumption will be given “exogenously”—consumption will equal labor income—and that the rich agent can be studied in isolation. Thus, the rich agent determines the interest rate and all the asset prices.

An important aspect of our approach is our emphasis on “rationality”: following recent work by Faruk Gul and Wolfgang Pesendorfer (2000a and 2000b), we have learned that behavior with “urges” can be rationalized within the standard economic theory paradigm. These authors have demonstrated that it is not necessary to interpret “preference reversals”—as doc-

¹Other limited-participation studies can be found in Basak and Cuoco (1998) and Dai (2001).

umented in the experimental literature—as expressions of irrationality or of “time inconsistency”. Instead, one can think of a consumer’s choices as depending not only on what they consume, but also on the set from which they make their choices. For example, a large set of possibilities may be strictly worse than a strict subset of the large set because the large set allows temptation—it contains tempting elements. The result of the temptation may either be that the consumer succumbs to the temptation or that he exercises self-control, but the consumer would be better off with the smaller set in either of these cases. Our contribution in the present paper is unimportant on a formal level; the framework we use is an uncertainty version of the model in our earlier paper (Krusell, Kuruşçu, and Smith (2001)) where we extend and specialize Gul and Pesendorfer’s work. This framework is attractive because of its ease of interpretation and its tractability. Moreover it is formally simply a generalization of the standard Lucas (1978) and Mehra and Prescott (1985) asset pricing model.

Our results are hard to judge quantitatively since we do not have any independent information regarding the strength or nature of the possible savings urges among investors. Clearly, if these urges are strong, the equilibrium risk-free rate will be very low. If we only introduce a savings urge, both the risk-free rate and the equity return will fall, so an increase in the equity premium cannot be obtained without other changes in the model. The natural way to obtain a high risk premium and a low risk-free rate, therefore, is to assume a high risk aversion and a strong savings urge. We show that this is possible in our model. In addition, our departure from the representative-agent setup itself can help in matching asset price moments. With more wealth concentration, the dividend risk is borne by a small group, and these investors have substantially more variable income growth than the average agent. Moreover, their labor income is positively but not perfectly correlated with the dividend risk, and they also receive a more than proportional share of total labor income. Under these conditions, we show that much less extreme values of risk aversion may be necessary in order to obtain a substantial equity premium, and that the variability of the risk-free rate need not be high at all.

Our focus on discounting is related to the argument in Kocherlakota (1990). He argues that a negative rate of discount can help solve the risk premium. What we offer is an interpretation of Kocherlakota’s point: it is possible to accept a negative rate of discount without having to accept a preference for consumption at any future date over current consumption (in his

paper, utils t periods from now are δ^t times more valuable than utils now, and since $\delta > 1$ this amount increases without bound). In other words, our deviation from constant discounting allows a distinction that Kocherlakota could not make. We make a similar point to the one in this paper in our Krusell, Kuruşçu, and Smith (2000) piece. There, we study a “Laibson model” (see, e.g., Laibson (1997)) and extend it to also allow Epstein-Zin (1989) preferences, allowing to disentangle risk aversion from intertemporal elasticity of substitution. Luttmer and Mariotti (2000) also studies asset pricing in the Laibson model (in much more depth than in our earlier paper); they, however, focus exclusively on consumption urges and on how the stochastic properties of the pricing kernel changes as a function of this urge. Finally, Gul and Pesendorfer themselves suggest self-control problems as a possible clue to asset pricing puzzles by suggesting that investors are “tempted to take on risk”—their temptation is to be risk-lovers (see Gul and Pesendorfer (2000b)).

In Section 2 we illustrate how the basic temptation model works for a single agent living in a two-period world without uncertainty. That model can then be extended to more periods, uncertainty, and more agents. Section 3 describes the full model. Section 4 contains our quantitative results and Section 5 concludes.

2 A 2-Period Illustration

Our 2-period model is now described. It is designed to illustrate the workings of the model. We will then briefly mention the effect of adding periods before we proceed to the infinite-horizon model.

A consumer in the economy values consumption today (c_1) and tomorrow (c_2). To express the idea of temptation, however, it is actually necessary to think of preferences over a different domain than simply pairs of period 1 and period 2 consumption levels: we need to define preferences over *sets* of feasible intertemporal consumption bundles. The reason why sets are needed is, loosely speaking, that temptation has to do with not just what a consumer chooses, but what he could have chosen. Thus, in a utility function representation of preferences, the agent’s utility function is defined over sets of those (c_1, c_2) pairs that the agent has available to him; in abstract, if A is an arbitrary set of consumption goods from which the consumer can make a

choice, his utility is given by

$$w(A) = \max_{(c_1, c_2) \in A} \{\tilde{u}(c_1, c_2) + \tilde{v}(c_1, c_2)\} - \max_{(\tilde{c}_1, \tilde{c}_2) \in A} \{\tilde{v}(\tilde{c}_1, \tilde{c}_2)\}.$$

That is, the agent's preferences can be described by two functions whose domain is pairs of consumption in periods 1 and 2: \tilde{u} and \tilde{v} . The former is referred to as the “commitment utility function” and the latter the “temptation utility function”. To see why these terms are used, notice that if the set A is a singleton (\bar{c}_1, \bar{c}_2) , we have $w(A) = \tilde{u}(\bar{c}_1, \bar{c}_2)$: if the agent has no choice—has committed to consumption—his preferences are given by \tilde{u} . If, on the other hand, A is not a singleton, then the utility is influenced by both functions.

The agent's choice out of the given set is that given by the argument that maximizes $\tilde{u} + \tilde{v}$. The solution to the second maximization problem—with \tilde{v} being the objective—has no interpretation in terms of observed choice. The second maximization problem, however serves to emphasize a “disutility of self-control”: $\tilde{v}(c_1, c_2) - \max_{(\tilde{c}_1, \tilde{c}_2) \in A} \{\tilde{v}(\tilde{c}_1, \tilde{c}_2)\}$ has to be nonpositive, and it is negative whenever \tilde{v} dictates a different choice than \tilde{u} .

For a brief (but not self-contained) background, Gul and Pesendorfer consider axioms over sets designed to (i) include the standard case without temptation but (ii) allow in addition a narrow set of circumstances that can be described by the idea of temptation and self-control. For this set of axioms, they provide the utility function characterization above: an agent satisfying the axioms can be uniquely identified with a pair of functions \tilde{u} and \tilde{v} . From the point of view of applications, the important point here is that these functions are independent of the particular choice problem the agent is faced with. Aside from basic axioms such as completeness, transitivity, and continuity, the additional key axiom is set betweenness: for any two sets A and B where A is preferred to B , it must also be that A is preferred to the union of A and B , and that the union of A and B is preferred to B : the union is “in between”. When A is strictly preferred to the union, we have a case of a “preference for commitment”: you strictly prefer the smaller set. A standard consumer (i.e., one who does not display a preference for commitment) would in contrast always be indifferent between A and $A \cup B$ whenever A is preferred to B . Moreover, when there is a preference for commitment, the agent “succumbs” to temptation if he is also indifferent between $A \cup B$ and B ; this is simply because when B is available, he would choose an element in B . On the other hand, the agent may not succumb

to temptation but rather exercise self-control: he might strictly prefer $A \cup B$ over B , because he chooses an element in A when facing the set $A \cup B$.

2.1 Our parameterization

We specify the functions \tilde{u} and \tilde{v} as follows:

$$\tilde{u}(c_1, c_2) = u(c_1) + \delta u(c_2)$$

and

$$\tilde{v}(c_1, c_2) = \gamma (u(c_1) + \beta \delta u(c_2)),$$

where u has the usual properties. For a standard consumer—one without a preference for commitment in any choice situation—we have that $\beta = 1$. When $\beta < 1$, the temptation function gives a stronger preference for present consumption. The parameter γ allows us to regulate the strength of the temptation; as it increases toward infinity, the agent moves closer to succumbing to temptation and let his actual choice be governed by \tilde{v} alone. As we shall see below, the use of the same functional form for \tilde{u} as for \tilde{v} is key for solving the model with ease. In addition, it also allows for a balanced growth equilibrium: when consumption grows at a constant rate, all contingent-claims prices will be constant.




Figure 1

Figure 1 illustrates the “urge to consume” for the case where the set from which the consumer chooses is a standard budget set. Here, the commitment

utility function would indicate that the best point in the budget set is where its indifference curve intersects the budget line. However, the temptation utility at this point has a slope greater than the interest rate. The maximization of $\tilde{u} + \tilde{v}$ would therefore lead to a higher level of consumption in period 1.

Figure 2 depicts the case we are primarily interested in here: it considers the possibility that $\beta > 1$. Now, the consumer will end up with a choice of consumption in period 1 below that maximizing \tilde{u} .




Figure 2

Formally, the utility of a typical consumer in this choice situation, then, is:

$$\max_{c_1, c_2} \{ \tilde{u}(c_1, c_2) + \tilde{v}(c_1, c_2) \} - \max_{\tilde{c}_1, \tilde{c}_2} \tilde{v}(\tilde{c}_1, \tilde{c}_2),$$

where both maximizations are subject to the budget constraint.

In this two-period problem, the “temptation” part of the problem (i.e., the second maximization problem in the objective function) plays no role in determining the consumer’s actions in period 1. As we describe just below, this is not true when the horizon is longer than two periods: then, the consumer’s temptation in future periods will influence the choices in the earlier periods, since the savings decisions earlier on influence the sizes of the budget sets later on, thus affecting the extent of temptation, with the associated disutility of self-control, in the later periods.

The consumer's intertemporal first-order condition is:

$$\frac{1 + \gamma}{(1 + \beta\gamma)\delta} \frac{u'(c_1)}{u'(c_2)} = 1 + r,$$

where r is the interest rate. It is straightforward to see that the intertemporal consumption allocation (which, in effect, maximizes $\tilde{u} + \tilde{v}$) represents a compromise between maximizing $\tilde{u} = u(c_1) + \delta u(c_2)$ and maximizing $\tilde{v} = u(c_1) + \beta\delta u(c_2)$. In the former case, the first-order condition is:

$$\frac{1}{\delta} \frac{u'(c_1)}{u'(c_2)} = 1 + r,$$

whereas in the latter case, the first-order condition is:

$$\frac{1}{\beta\delta} \frac{u'(c_1)}{u'(c_2)} = 1 + r.$$

Since

$$\frac{1}{\beta\delta} \leq \frac{1 + \gamma}{(1 + \beta\gamma)\delta} \leq \frac{1}{\delta}$$

in the case we are interested in, the consumer's consumption allocation is tilted towards the future relative to maximizing $u(c_1) + \delta u(c_2)$ and is tilted towards the present relative to maximizing the temptation function $u(c_1) + \beta\delta u(c_2)$.

2.1.1 Equilibrium

Equilibrium prices in a deterministic, representative-agent 2-period model are obtained simply by setting consumption equal to endowment: the interest rate, or one over the price of a bond, will then satisfy

$$1 + r = \frac{1}{q} = \frac{1 + \gamma}{(1 + \beta\gamma)\delta} \frac{u'(y_1)}{u'(y_2)},$$

where y_t is total endowment in period t . We see that, in the 2-period model, the equilibrium prices are proportional to the effective discount rate, $\frac{1+\gamma}{(1+\beta\gamma)\delta}$. With a higher β (and, given a $\beta > 1$, a higher γ) interest rates fall. This is the main mechanism leading to a lower risk-free rate also in the stochastic, multi-period model, although the situation is more complicated there. Observe also that β and δ are observationally equivalent in the 2-period model (for a lower risk-free rate, an alternative to a higher β is of course a higher δ).

2.1.2 Uncertainty

Uncertainty in the 2-period model is easy to analyze, too. First, however, one needs to be clear on what constitutes temptation in this case. We will simply generalize from the deterministic case, which means that we assume that the attitudes toward risk are not fundamentally different: the temptation function \tilde{v} embodies the same curvature assumptions as the commitment function \tilde{u} . Here, Gul and Pesendorfer (2000b) make a different choice—they consider \tilde{v} to be *convex* and thereby consider a temptation function that embodies risk-loving.

With uncertainty in period 2, and a contingent claim for every state of the world in period 2, we arrive at simple expressions for the prices of these claims:

$$q_i = \pi_i \frac{1 + \beta\gamma}{1 + \gamma} \delta \frac{u'(y_{2i})}{u'(y_1)}.$$

The risk-free rate, similarly, is given by $1/\sum_i q_i$, so that

$$\frac{1}{1+r} = \frac{1 + \beta\gamma}{1 + \gamma} \delta \sum_i \pi_i \frac{u'(y_{2i})}{u'(y_1)}.$$

The price of a tree with dividend d_i in state i is thus

$$p = \frac{1 + \beta\gamma}{1 + \gamma} \delta \sum_i \pi_i \frac{u'(y_{2i})}{u'(y_1)} d_i.$$

The ratio of the expected gross return on equity, $\sum_i \pi_i d_i/p$, to the gross risk-free rate, is therefore

$$\sum_i \pi_i d_i \frac{\sum_i \pi_i u'(y_{2i})}{\sum_i \pi_i u'(y_{2i}) d_i}.$$

We see that the (gross) equity premium—expressed in this proportional form—does not depend on the discount rates; the key determinant of the equity premium is how the dividend correlates with the marginal utility of consumption. This can be seen by rewriting the expression as

$$E[d] \frac{E(\text{MU})}{E(d \cdot \text{MU})} = \frac{E(\text{MU})}{E\left(\text{MU} \cdot \frac{d}{E[d]}\right)} = \frac{1}{\text{Cov}\left(\frac{\text{MU}}{E[\text{MU}]}, \frac{d}{E[d]}\right) + 1}.$$

As in the standard model, the equity premium is positive provided that marginal utility covaries negatively with the dividend.

The main insight, then, of this 2-period model with uncertainty is that urges to save or consume (as captured by the parameter β) affect the risk-free rate but not the equity premium. This insight carries over, in large measure, to the infinite-horizon model that we describe in Section 3. Although β does have small effects on the equity premium in this model, the main effect of changes in β is to change the levels of the rates of return but not their spread.

2.2 More than 2 periods

With more than two periods, we first need to specify what an agent is tempted by. We will assume that temptation only involves the immediate present. That is, we will assume that \tilde{u} and \tilde{v} agree on any future variables: holding current consumption constant, these functions are identical. Formally, this means that if A and B are sets of intertemporal consumption bundles and if A and B have identical possibilities for current consumption, then either $A \cup B$ is indifferent to A or it is indifferent to B .

These assumptions allow us to solve the problem backwards. One first analyzes the last two periods, $T-1$ and T , in a manner parallel to that in the previous subsection. In the context of budget sets again, this gives rise to a utility over sets $w_{T-1}(\omega_{T-1})$, where ω_{T-1} is incoming wealth in period $T-1$. The third period from the last, the utility of any ω_{T-2} can then be expressed using the indirect utility from the previous step. Our assumption will now be that the commitment utility attaches a weight of δ on $w_{T-1}(\omega_{T-1})$, whereas the temptation utility attaches a weight $\beta\delta$. This is the natural generalization of the two-period case, and the formulation we have adopted in our earlier work (Krusell, Kuruşçu, and Smith (2001)). Our earlier paper also shows that this formulation—a “quasi-geometric” temptation—turns out to work as a generalization of sorts of the Laibson model.

To look in some more detail at the multi-period model, consider the three-period case: in period 2, the consumer will be given some initial wealth that was determined in period 1 and chooses over consumption levels in periods 2 and 3 as in the above subsection. However, how is the choice in period 1 determined? Here, since we are considering a choice of period-2 budget sets, the agent in period 1 must consider not only the consumption outcomes in periods 2 and 3 but also the possible disutility of temptation in period 2. (Period 3, of course, does not admit such a disutility, since the choice set in period 3 is a singleton.) It is straightforward to show that the first-order

condition for savings in period 1 now becomes

$$u'(c_1) = \delta \frac{1 + \beta\gamma}{1 + \gamma} (1 + r) [u'(c_2) + \gamma (u'(c_2) - u'(\tilde{c}_2))],$$

where \tilde{c}_2 as before refers to temptation consumption in period 2. This expression looks like a standard Euler equation except in two places: (i) the discount factor is different; and (ii) there is an added term $\gamma(u'_{t+1} - \tilde{u}'_{t+1})$. The discount factor, which is the same as in the 2-period model, is greater than δ . The term $\gamma(u'_2 - \tilde{u}'_2)$ is the derivative of the disutility/cost of self-control, $\gamma(u_2 - \tilde{u}_2)$, with respect to wealth. This term is negative, since temptation consumption is lower than actual consumption when $\beta > 1$ and the utility function is strictly concave. The interpretation of these equations is that the marginal benefit from wealth tomorrow falls short of u'_2 , because the self-control cost gets larger as wealth increases in this model!

Similarly, there will be an Euler equation for the maximization of temptation utility in period 1:

$$u'(\tilde{c}_1) = \beta\delta(1 + r) [u'(\tilde{c}_2) + \gamma (u'(\tilde{c}_2) - u'(\tilde{\tilde{c}}_2))],$$

where \tilde{c}_2 is period 1 temptation consumption, \tilde{c}_2 is period 2 actual consumption if you succumb in period 1, and $\tilde{\tilde{c}}_2$ is period 2 temptation consumption if you succumb in period 1.

Equilibrium here is significantly harder to solve for. Not only does temptation influence actual decisions, but one also cannot easily find out what equilibrium temptation consumption is, even in a deterministic endowment economy—equilibrium actual consumption equals endowments but temptation consumption has the nature of a “deviation” from equilibrium. With uncertainty, the risk premium will now involve the covariance between dividends and $(1 + \gamma)\text{MU} - \gamma\widetilde{\text{MU}}$, so the stochastic process for temptation consumption becomes relevant. In general, there will also not be observational equivalence with the the standard model. We will now move on to the infinite-horizon case, where uncertainty will indeed play a central role, as will borrowing constraints.

3 The Infinite-Horizon Asset Pricing Model

3.1 Primitives

Consider now an infinite-horizon setup with two agents: the big guy (B) and the little guy (L). The agents have clones: there is a measure θ of big guys and a measure $1 - \theta$ of little guys.

As in the Lucas asset pricing model, there is no physical capital accumulation nor production: total consumption will equal total endowments. The aggregate endowment process has two parts: there is a “dividend from a tree”, d , as well as “labor income”, l . As in Mehra and Prescott (1985), there is dividend growth, and labor income will grow too. Total income, y , grows stochastically at one of two possible rates. The dividend is specified as a fraction η of total output: $d = \eta y$, with an implied $l = (1 - \eta)y$. The dividend share of output is also stochastic but has a four-state support. Therefore, our stochastic process for endowments has four possible states. We denote these by 1 through 4; states 1 and 3 have low and states 2 and 4 high output growth, and the dividend/output ratio is increasing from state 1 through state 4. Associated to each state i are four transition probabilities π_{ij} for $j = 1, 2, 3, 4$.

We will use an asset structure with contingent claims which will each be zero in net supply (see the next section); therefore, the endowment process for each agent can be described by how much of the dividend and how much of the labor income he receives. We will assume that the little guy gets no dividend income and that, in aggregate, a fraction ψ of labor income goes to the big guys and a fraction $1 - \psi$ to the little guys. In order to express the idea that the big guy’s income (and consumption) varies more like the dividend process itself we will make ψ stochastic. More specifically, it will take on one of four possible values: one for each of the above states j . Each little guy will therefore consume his endowment $y^L = l(1 - \psi)/(1 - \theta)$ in equilibrium and each big guy will consume $y^B = d/\theta + l\psi/\theta$. Our calibration below will make ψ much larger than θ (to match income inequality) and positively correlated with η .

We also should describe the preferences here; essentially, they can be thought of as the extension of the finite-period case to an infinite horizon. However, we postpone this task since it is much easier to describe the preferences in the context of a specific choice problem, which requires a description of the asset structure.

3.2 Assets and asset prices

The asset structure is one with contingent claims: there are four such claims for each date and state, each delivering one unit of consumption in each of the consecutive four states. The price of a claim in state i delivering 1 unit of consumption in state j next period is denoted q_{ij} . The contingent claims allows us to easily compute the prices of stocks and bonds. The stock prices, p_i for states $i = 1, 2, 3, 4$, will be given by the solution the four equations

$$p_i = q_{i1}(d_1 + p_1) + q_{i2}(d_2 + p_2) + q_{i3}(d_3 + p_3) + q_{i4}(d_4 + p_4)$$

for $i = 1, 2, 3, 4$. Clearly, the contingent-claims prices will typically uniquely define the stock prices. Similarly, the bond prices, q_i , are given by

$$q_i = q_{i1} + q_{i2} + q_{i3} + q_{i4}$$

for $i = 1, 2, 3, 4$.

The agents have no endowments of assets; the market-clearing level of each contingent claim is zero. The agents' endowments simply consist of the dividend and labor income processes. We assume that the big guy's endowment equals the dividend plus the labor income, and that the little guy's endowment is the labor income only.²

3.3 Borrowing constraints

We assume that agents face borrowing constraints. The brutal honest truth about these constraints is that we impose them in order to facilitate the solution of the model. With our approach, as we have already indicated, one of the two kinds of agents will be de facto more impatient than the other, and with a borrowing constraint the more impatient of the two—the little guy—will end up always consuming his labor endowment. This modelling approach really is a shortcut for something which we believe is more reasonable—to have an incomplete set of asset markets. Krusell and Smith (1997 and 1998) specify such a model; there, no agent faces a binding

²We could have chosen another endowment pattern. However, what is important for our quantitative results is how much of total income we let the little guy consume. How much he is allowed to consume is a function both of his endowment and of the borrowing constraint he is faced with: for any assumption on the endowment level, there is a set of borrowing constraints that let us vary his consumption level, and vice versa.

constraint eternally, since discount factors are assumed to be stochastic, and solution of that model is therefore much more involved. However, the paper shows that small differences in discount factors lead to large differences in wealth holdings and savings behavior, and our present approach can merely be viewed as a simple way to mimic these results.

The implementation of borrowing constraints with contingent claims requires a specification of several constraints. The simplest way to do this is to require that all holdings of contingent claims be nonnegative; this is also what we do.³ In the formulation of the big guy's problem below, we will for simplicity omit the borrowing constraints; however, once a solution is found one has to show that the solution to the agent's problem indeed is interior.

3.4 The big guy's problem

We now specify the big guy's problem. We use recursive methods to get at the limit of the finite-horizon problems discussed above.⁴ It is given by the following recursive functional equation:

$$\begin{aligned}
 W_i(\omega, y) = & \tag{1} \\
 (1 + \gamma) \max_{\{s'_j\}_{j=1}^4} & \left\{ u \left(\omega - \sum_{j=1}^4 q_{ij} s'_j \right) + \delta \left(\frac{1 + \beta\gamma}{1 + \gamma} \right) \sum_{j=1}^4 \pi_{ij} W_j(\omega'_j, y'_j) \right\} - \\
 \gamma \max_{\{\tilde{s}'_j\}_{j=1}^4} & \left\{ u \left(\omega - \sum_{j=1}^4 q_{ij} \tilde{s}'_j \right) + \beta\delta \sum_{j=1}^4 \pi_{ij} W_j(\tilde{\omega}'_j, y'_j) \right\},
 \end{aligned}$$

where

$$\begin{aligned}
 \omega'_j &= s'_j + y'_j, \\
 \tilde{\omega}'_j &= \tilde{s}'_j + y'_j, \\
 y'_j &= g_{ij} y,
 \end{aligned}$$

³There are alternatives as well, including constraints on total wealth, margin constraints, and so on.

⁴Axiomatization leading to the recursive version of the Gul-Pesendorfer preferences with quasi-geometric temptation that we consider here is not fully established. See Krusell, Kuruşçu, and Smith (2001) for details.

and g_{ij} is the growth rate of the big guy's endowment, given that the economy is in state i today and in state j tomorrow. (To simplify notation, we are using y to denote y^B , the big guy's endowment.) This problem delivers decision rules for actual behavior, $\{s'_{ij}(\omega, y)\}_{j=1}^4$, and temptation behavior, $\{\tilde{s}'_{ij}(\omega, y)\}_{j=1}^4$.

3.5 Finding an equilibrium

Algorithmically, it is clear how to find an equilibrium for our economy. We first solve the problem of the big guy in isolation to obtain market-clearing prices and behavior that is optimal given these prices. We then verify that, given these prices, it is indeed optimal for the little guy to have a binding borrowing constraint; this is established using the appropriate Euler inequalities for all four states.

To find the market-clearing prices in this economy is significantly more complicated, unfortunately, than in the Mehra and Prescott economy. In the latter, contingent claims prices can be obtained immediately from the first-order conditions for each contingent claim, since consumption has to equal endowments. In the present economy too, actual consumption has to equal endowment, but it is also necessary to find the temptation consumption levels, since these appear in the first-order conditions. Finding the temptation consumption levels is what is hard. To see this, let us study the first-order conditions in some detail.

3.5.1 The Euler equations

There are first-order conditions both for actual decision rules and for the temptation decision rules, as we saw in the deterministic case above. In the case of uncertainty, the conditions read as follows:

$$q_{ij}u' \left(\omega - \sum_{l=1}^4 q_{il}s'_{il}(\omega, y) \right) = \delta \frac{1 + \gamma\beta}{1 + \gamma} \pi_{ij} \cdot$$

$$\left[(1 + \gamma)u' \left(s'_{ij}(\omega, y) + y'_j - \sum_{l=1}^4 q_{jl}s'_{jl}(\omega'_j, y'_j) \right) - \right.$$

$$\left. \gamma u' \left(s'_{ij}(\omega, y) + y'_j - \sum_{l=1}^4 q_{jl}\tilde{s}'_{jl}(\omega'_j, y'_j) \right) \right]$$

for actual behavior and

$$q_{ij}u' \left(\omega - \sum_{l=1}^4 q_{il}\tilde{s}'_{il}(\omega, y) \right) = \delta\beta\pi_{ij} \cdot$$

$$\left[(1 + \gamma)u' \left(\tilde{s}'_{ij}(\omega, y) + y'_j - \sum_{l=1}^4 q_{jl}s'_{jl}(\tilde{\omega}'_j, y'_j) \right) - \right.$$

$$\left. \gamma u' \left(\tilde{s}'_{ij}(\omega, y) + y'_j - \sum_{l=1}^4 q_{jl}\tilde{s}'_{jl}(\tilde{\omega}'_j, y'_j) \right) \right]$$

for temptation behavior, where

$$\omega'_j = s'_{ij}(\omega, y) + y'_j$$

and

$$\tilde{\omega}'_j = \tilde{s}'_{ij}(\omega, y) + y'_j.$$

Each of these equations must hold for all ω , i , and j .

These equations look complicated because they are expressed as functional equations in the decision rules s' and \tilde{s}' : these decision rules are evaluated explicitly at the relevant values. For a full understanding of the conditions describing fully rational behavior, the reader is invited to carefully distinguish how these two functions appear in the equations.

On a conceptual level, what is important about the first-order equations is that they have to hold for all ω and not just for equilibrium wealth (y^B). In the temptation behavior, the consumer is choosing to save a different amount than that actually saved. A consequence of such savings would be that next period's wealth would be different than next period's actual wealth and, as a result, all subsequent periods would produce non-equilibrium wealth levels. To verify optimality of temptation behavior, therefore, it is necessary to know what consumption levels result from each possible wealth level: it is necessary to solve for entire (actual and temptation) decision rules. This is what makes the determination of asset prices in this economy a level more difficult than in the Mehra and Prescott economy.

3.5.2 Guessing at a specific solution

Fortunately, under the assumption that u is an isoelastic function, we can use a guess-and-verify approach to find the solution to the big guy's problem.

In particular, our guess will be not only that all decision rules are linear in wealth, but also that the allocation of risk is not affected by the temptation: the ratio of consumption levels in different states tomorrow are the same under actual as under temptation behavior.

Normally, the guess-and-verify approach would only require using the Euler equation: insert the guess for the decision rule into the Euler equation and show that it is possible to find decision rule coefficients that make the Euler equation hold for all values of the state variables. In this problem, however, it is helpful to study the functional equation itself. To this end, we now return to the functional equation and at the end, we will revisit the Euler equations.

We explain our solution to the functional equation in a series of steps. At the end of this section, we also discuss the solution to the little guy's problem. In the next section, we outline our computational algorithm, the details of which are relegated to the Appendix. A reader who is not interesting in the details of how we compute equilibrium can safely skip to Section 4, where we discuss our quantitative findings.

Step 1: Conjecture functional forms

As a first step, we display the conjectured functional forms for the value function and decision rules. To show the generality of our approach, let there be N exogenous states and N^2 state-contingent assets (in our problem, $N = 4$). Specifically, given a set of prices q_{ij} , $i = 1, \dots, N$, $j = 1, \dots, N$, we show below that the value function and decision rules have the following functional forms when $u(c) = c^\alpha/\alpha$:

$$\begin{aligned} W_i(\omega, y) &= A_i (D_i y + \omega)^\alpha, \quad i = 1, \dots, N, \\ s'_{ij}(\omega, y) &= a_{ij}y + b_{ij}\omega, \quad i = 1, \dots, N, \quad j = 1, \dots, N, \\ \tilde{s}'_{ij}(\omega, y) &= \tilde{a}_{ij}y + \tilde{b}_{ij}\omega \quad i = 1, \dots, N, \quad j = 1, \dots, N, \end{aligned}$$

where the A_i 's, D_i 's, a_{ij} 's, b_{ij} 's, \tilde{a}_{ij} 's, and \tilde{b}_{ij} 's are unknown coefficients for which we must solve. In total, there are $2N + 4N^2$ of these coefficients. In the rest of this section, we describe a set of $2N + 4N^2$ equations that these coefficients must satisfy.

Step 2: Normalize by income and state the functional equation

Next, we eliminate y from the functional equation that defines the big guy's problem. In particular, given the functional forms above, the common term y^α can be cancelled from both sides of the functional equation. Defining $\hat{\omega} = \omega/y$, $\hat{s}'_{ij} = s'_{ij}/y$, and $\hat{\tilde{s}}'_{ij} = \tilde{s}'_{ij}/y$, the functional equation can then be written:

$$A_i (D_i + \hat{\omega})^\alpha = (1 + \gamma) \cdot \quad (2)$$

$$\max_{\{\hat{s}'_{ij}\}_{j=1}^4} \left\{ \alpha^{-1} \left(\hat{\omega} - \sum_{j=1}^N q_{ij} \hat{s}'_{ij} \right)^\alpha + \delta \left(\frac{1 + \beta\gamma}{1 + \gamma} \right) \sum_{j=1}^N A_j \pi_{ij} (D_j + \hat{w}'_j)^\alpha g_{ij}^\alpha \right\} -$$

$$\gamma \max_{\{\hat{\tilde{s}}'_{ij}\}_{j=1}^4} \left\{ \alpha^{-1} \left(\hat{\omega} - \sum_{j=1}^N q_{ij} \hat{\tilde{s}}'_{ij} \right)^\alpha + \beta\delta \sum_{j=1}^N A_j \pi_{ij} (D_j + \hat{w}'_j)^\alpha g_{ij}^\alpha \right\},$$

where

$$\hat{w}'_j = \hat{s}'_{ij}/g_{ij} + 1$$

and

$$\hat{\tilde{w}}'_j = \hat{\tilde{s}}'_{ij}/g_{ij} + 1.$$

Note that, given the conjectured functional forms above, \hat{s}'_{ij} and $\hat{\tilde{s}}'_{ij}$ are each affine functions of $\hat{\omega}$.

Given the guess for $\hat{s}'_{ij}(\hat{\omega})$, actual consumption (divided by y) in state i can be written:

$$\hat{c}_i(\hat{\omega}) = \hat{\omega} - \sum_{j=1}^N q_{ij} \hat{s}'_{ij}(\hat{\omega}) = \left(1 - \sum_{j=1}^N q_{ij} b_{ij} \right) \left(\frac{-\sum_{j=1}^N q_{ij} a_{ij}}{1 - \sum_{j=1}^N q_{ij} b_{ij}} + \hat{\omega} \right). \quad (3)$$

Similarly, given the guess for $\hat{\tilde{s}}'_{ij}(\hat{\omega})$, temptation consumption (divided by y) in state i can be written:

$$\hat{\tilde{c}}_i(\hat{\omega}) = \hat{\omega} - \sum_{j=1}^N q_{ij} \hat{\tilde{s}}'_{ij}(\hat{\omega}) = \left(1 - \sum_{j=1}^N q_{ij} \tilde{b}_{ij} \right) \left(\frac{-\sum_{j=1}^N q_{ij} \tilde{a}_{ij}}{1 - \sum_{j=1}^N q_{ij} \tilde{b}_{ij}} + \hat{\omega} \right). \quad (4)$$

We use these expressions extensively below.

Step 3: Verify the functional form of the value function

Examining the right-hand side of the functional equation (2), one can see that there are $2(N+1)$ terms involving $\hat{\omega}$ ($N+1$ terms in the ‘‘actual’’ part of

the functional equation and $N + 1$ analogous terms in the “temptation” part of the functional equation). Two of these terms are the ones above for actual and temptation consumption (raised to the power of α). The remaining $2N$ of these terms either take the form

$$\left(D_j + 1 + \hat{s}'_{ij}/g_{ij}\right)^\alpha = (b_{ij}/g_{ij})^\alpha \left(\frac{D_j + 1 + a_{ij}/g_{ij}}{b_{ij}/g_{ij}} + \hat{w}\right)^\alpha, \quad j = 1, \dots, N,$$

or the form

$$\left(D_j + 1 + \tilde{s}'_{ij}/g_{ij}\right)^\alpha = (\tilde{b}_{ij}/g_{ij})^\alpha \left(\frac{D_j + 1 + \tilde{a}_{ij}/g_{ij}}{\tilde{b}_{ij}/g_{ij}} + \hat{w}\right)^\alpha, \quad j = 1, \dots, N.$$

Finally, the left-hand side of the functional equation also contains a term involving ω : $(D_i + \hat{\omega})^\alpha$.

If the conjectured functional forms are correct, then they must solve the functional equation for all values of $\hat{\omega}$. In other words, the terms involving $\hat{\omega}$ must cancel from both sides of the functional equation. This can only happen if all the affine expressions involving $\hat{\omega}$ (that are all raised to a power α) have the same relation between slope and intercept. This produces the following restrictions, for each $i = 1, \dots, N$:

$$\frac{-\sum_{j=1}^N q_{ij} a_{ij}}{1 - \sum_{j=1}^N q_{ij} b_{ij}} = \frac{D_j + 1 + a_{ij}/g_{ij}}{b_{ij}/g_{ij}}, \quad j = 1, \dots, N, \quad (5)$$

and

$$\frac{-\sum_{j=1}^N q_{ij} \tilde{a}_{ij}}{1 - \sum_{j=1}^N q_{ij} \tilde{b}_{ij}} = \frac{D_j + 1 + \tilde{a}_{ij}/g_{ij}}{\tilde{b}_{ij}/g_{ij}}, \quad j = 1, \dots, N. \quad (6)$$

Finally, for each $i = 1, \dots, N$, we also require that

$$D_i = \frac{-\sum_{j=1}^N q_{ij} a_{ij}}{1 - \sum_{j=1}^N q_{ij} b_{ij}} = \frac{-\sum_{j=1}^N q_{ij} \tilde{a}_{ij}}{1 - \sum_{j=1}^N q_{ij} \tilde{b}_{ij}}. \quad (7)$$

The set of equations (5), (6), and (7) defines $2N + 2N^2$ restrictions that the unknown coefficients must satisfy. We now turn to the first-order conditions, which define another set of restrictions that the unknown coefficients must satisfy.

Step 4: Verify the functional forms of the decision rules

The first-order conditions governing optimal (actual and temptation) asset choices deliver an additional $2N^2$ restrictions on the coefficients. Given the guess for the value function, the first-order conditions for actual behavior take the following form for $i = 1, \dots, N$ and $j = 1, \dots, N$:

$$q_{ij} \left(\hat{w} - \sum_{j=1}^N q_{ij} \hat{s}'_{ij}(\hat{w}) \right)^{\alpha-1} = \left(\frac{1 + \beta\gamma}{1 + \gamma} \right) \pi_{ij} \delta \alpha A_j \left(D_j + 1 + \hat{s}'_{ij}(\hat{w}) / g_{ij} \right)^{\alpha-1} g_{ij}^{\alpha-1}.$$

Substituting the guess for $\hat{s}'_{ij}(\hat{w})$ and using equation (5), the terms involving \hat{w} cancel, thereby verifying that the conjectured functional form for the actual decision rules is correct.

Step 5: Collect all the restrictions on unknown coefficients

After eliminating \hat{w} , the N^2 first-order conditions for actual behavior become:

$$q_{ij} \left(1 - \sum_{j=1}^N q_{ij} b_{ij} \right)^{\alpha-1} = \left(\frac{1 + \beta\gamma}{1 + \gamma} \right) \pi_{ij} \delta \alpha A_j b_{ij}^{\alpha-1}. \quad (8)$$

Similarly, the N^2 first-order conditions for temptation behavior reduce to:

$$q_{ij} \left(1 - \sum_{j=1}^N q_{ij} \tilde{b}_{ij} \right)^{\alpha-1} = \pi_{ij} \beta \delta \alpha A_j \tilde{b}_{ij}^{\alpha-1}. \quad (9)$$

From before, we have restrictions in the form of equations (5)–(7). The final N restrictions that the unknown coefficients must satisfy come directly from the functional equation. After cancelling terms involving \hat{w} , thanks to equations (5)–(7), the functional equation, for each $i = 1, \dots, N$, can be written:

$$A_i = (1 + \gamma) \left\{ \alpha^{-1} \left(1 - \sum_{j=1}^N q_{ij} b_{ij} \right)^{\alpha} + \left(\frac{1 + \beta\gamma}{1 + \gamma} \right) \delta \sum_{j=1}^N \pi_{ij} A_j b_{ij}^{\alpha} \right\} - \gamma \left\{ \alpha^{-1} \left(1 - \sum_{j=1}^N q_{ij} \tilde{b}_{ij} \right)^{\alpha} + \beta \delta \sum_{j=1}^N \pi_{ij} A_j \tilde{b}_{ij}^{\alpha} \right\}. \quad (10)$$

In sum, the unknown coefficients in the value function and in the actual and temptation decision rules must satisfy the $3N + 4N^2$ (independent) equations contained in (5)–(10). But recall that there are only $2N + 4N^2$ coefficients in the value function and decision rules! In the next step, we show that N of the restrictions described above are actually redundant.

Step 6: Show that N of the restrictions are redundant

It turns out the N of the equations in (5)–(7) are redundant. To see this, note that, for each i , equation (5) implies that:

$$q_{ij}g_{ij}(D_j + 1) = q_{ij}b_{ij} \left(\frac{-\sum_{j=1}^N q_{ij}a_{ij}}{1 - \sum_{j=1}^N q_{ij}b_{ij}} \right) - q_{ij}a_{ij}, \quad j = 1, \dots, N.$$

Summing these equations over j and simplifying the right-hand side, one obtains:

$$\sum_{j=1}^N q_{ij}g_{ij}(D_j + 1) = \frac{-\sum_{j=1}^N q_{ij}a_{ij}}{1 - \sum_{j=1}^N q_{ij}b_{ij}}.$$

By equation (7), for the given i ,

$$\sum_{j=1}^N q_{ij}g_{ij}(D_j + 1) = \frac{-\sum_{j=1}^N q_{ij}\tilde{a}_{ij}}{1 - \sum_{j=1}^N q_{ij}\tilde{b}_{ij}}.$$

Clearly, if this equation and the first $N - 1$ of the equations in (6) hold (for the given i), then the last equation in (6) also holds. Thus, for each $i = 1, \dots, N$, there is one redundant equation.

Finally, for the little guy the first-order conditions (Euler equations) must hold as inequalities: the left-hand side has to exceed the right-hand side when evaluated at $s' = 0$. We assume that the little guy has no urge to save (that his β is less than or equal to one) so that \tilde{s}' is less than or equal to s' : thus, if the borrowing constraint binds for actual behavior, it will for temptation behavior.⁵

⁵This point was made by Gul and Pesendorfer in their (2000b) paper: when faced with a binding borrowing constraint, an agent with an urge to consume does not suffer from a disutility of self-control, and can thus be viewed as observationally equivalent to one without an urge to consume.

3.5.3 Outline of the computational algorithm

Because equilibrium wealth $\hat{\omega} = 1$, market-clearing in each of the N^2 asset markets requires that

$$a_{ij} + b_{ij} = 0, \quad i = 1, \dots, N, \quad j = 1, \dots, N. \quad (11)$$

A direct approach to computing equilibrium would be to use a nonlinear equation solver to find the solution to equations (5)–(10) and (11). In total, this set of equations has $2N + 5N^2$ unknowns (N^2 prices in addition to $2N + 4N^2$ coefficients in the value function and decision rules). For our problem, in which $N = 4$, there are, therefore, 88 nonlinear equations in 88 unknowns. Because this system is large, we use instead a less direct but more manageable algorithm described below.

The basic idea of the algorithm is to fix a set of prices, solve the big guy’s problem, and then check whether markets clear. If they do not, update the prices using a nonlinear equation solver and continue iterating until convergence. Thus, the “outer loop” of our computational algorithm requires the solution of 16 equations, i.e., the ones given in equation (11), in the 16 prices.

The “inner loop” of our computational algorithm solves the big guy’s problem given a set of prices. Here, we exploit some additional structure in the problem to reduce the computational burden. In particular, dividing equation (8) by equation (9) for a given i and j yields a useful relationship between the slope coefficients in the actual and temptation decision rules:

$$\frac{1 - \sum_{j=1}^N q_{ij} b_{ij}}{1 - \sum_{j=1}^N q_{ij} \tilde{b}_{ij}} = \left(\frac{1 + \beta\gamma}{(1 + \gamma)\beta} \right)^{\frac{1}{\alpha-1}} \frac{b_{ij}}{\tilde{b}_{ij}}.$$

Note that the left-hand side of this equation is simply the ratio of actual consumption to temptation consumption in state i , which one can see by dividing the right-hand side of equation (3) by the right-hand side of equation (4) and using equation (7) to eliminate the terms involving $\hat{\omega}$. Moreover, because the ratio b_{ij}/\tilde{b}_{ij} depends only on i , this equation implies that

$$\frac{c_i(\hat{\omega})}{\tilde{c}_i(\hat{\omega})} = \frac{1 - \sum_{j=1}^N q_{ij} b_{ij}}{1 - \sum_{j=1}^N q_{ij} \tilde{b}_{ij}} = \left(\frac{1 + \beta\gamma}{(1 + \gamma)\beta} \right)^{\frac{1}{\alpha-1}} \frac{\sum_{j=1}^N q_{ij} b_{ij}}{\sum_{j=1}^N q_{ij} \tilde{b}_{ij}}. \quad (12)$$

Evidently, this consumption ratio depends only on i and not on $\hat{\omega}$; this fact is useful for computation. As discussed below, we exploit this equation to simplify the computation of the decision rule coefficients.

Our computational algorithm uses only the first-order conditions associated with the big guy's problem, obviating the need to compute the $2N$ coefficients in the value function. Because $N = 4$ in our problem, we must therefore compute $2N^2 = 32$ coefficients in the actual decision rules, as well as 32 coefficients in the temptation decision rules, given a set of prices. A typical first-order condition for actual behavior associated with the problem defined by equation (2) can be written as:

$$q_{ij}c^{\alpha-1} = \delta \left(\frac{1 + \beta\gamma}{1 + \gamma} \right) \pi_{ij}g_{ij}^{\alpha-1} \left((1 + \gamma) (c'_j)^{\alpha-1} - \gamma (\tilde{c}'_j)^{\alpha-1} \right),$$

where c is actual consumption in the current period, c'_j is actual consumption in the next period given that the economy is in state j , and \tilde{c}'_j is temptation consumption in the next period given that the economy is in state j . This first-order condition can be rearranged into a more useful form:

$$c = \left(\frac{\delta(1 + \beta\gamma)}{1 + \gamma} \right)^{\frac{1}{\alpha-1}} \left(\frac{\pi_{ij}}{q_{ij}} \right)^{\frac{1}{\alpha-1}} g_{ij} c'_j \left(1 + \gamma - \gamma \left(\frac{\tilde{c}'_j}{c'_j} \right)^{\alpha-1} \right)^{\frac{1}{\alpha-1}}. \quad (13)$$

This equation is central to our computational algorithm. Because temptation consumption enters this equation only in the form of a ratio to actual consumption, and because equation (12) shows how to compute this ratio given only the slope coefficients in the actual decision rules, we do not need to consider separately the first-order conditions for temptation consumption. In other words, thanks to equation (12), we need to solve only for the 32 coefficients in the actual decision rules, given a set of prices. The Appendix explains in detail how we use equation (13) to compute the actual decision rule coefficients given a set of prices, and how we update the prices to achieve market-clearing.

4 Quantitative Results

Although a key purpose of our model is to examine the role of the different parameters for asset prices—comparative-statics exercises—some parameters will be calibrated and remain unchanged throughout the analysis. We therefore first describe this calibration. Thereafter we show our main results.

4.1 Calibration

We took the dividend/labor income process from Heaton and Lucas (1996): they consider a four-state Markov process, which was matched to U.S. data.⁶ The transition probability matrix we thus use (it is derived from Table 1 in Heaton and Lucas’s paper) is

$$\begin{pmatrix} 0.5297 & 0.3024 & 0.1068 & 0.0611 \\ 0.4101 & 0.4675 & 0.0572 & 0.0652 \\ 0.0652 & 0.0572 & 0.4675 & 0.4101 \\ 0.0611 & 0.1068 & 0.3024 & 0.5297 \end{pmatrix}.$$

The (i, j) element in this matrix is the probability of moving from state i to state j . The four states are $(0.9904, 0.1402)$, $(1.047, 0.1437)$, $(0.9904, 0.1561)$, and $(1.047, 0.1599)$, where the first number in each ordered pair is the growth rate of total income (labor income plus dividends) and the second number in each ordered pair is the share of dividends in total income.⁷

As for the remaining parameters, we vary them away from a baseline configuration: $\delta = 0.99$, $\theta = 1$ (so that there is no little guy—the representative-agent version of the model), $\beta = 1$ and/or $\gamma = 0$ (the standard model), and $\alpha = 0$ (the coefficient of relative risk aversion being $1 - \alpha$).

In Sections 4.2.1, 4.2.2, and 4.2.3, we examine how changes in risk aversion (α), wealth concentration (θ), and the “urge to save” (β) affect equilibrium asset prices. In a reasonable calibration of the model, however, we should set wealth concentration, as well as income concentration, to match measures of concentration in the data. We do this toward the end of our parameter variations—in Section 4.2.4—where we fix θ and choose α and β so as to

⁶In particular, Heaton and Lucas (1996) obtain measures of annual aggregate labor income and dividends from the NIPA for the years 1947–1992. They weight the aggregate series by the total U.S. population and by the CPI in each year in order to obtain real per capita labor income and dividends. They fit a first-order vector autoregression to the growth rate of total income (labor income plus dividends) and the share of dividends in total income and then use the technique described in Tauchen and Hussey (1991) to approximate this vector autoregression using a four-state Markov chain. In the NIPA, the average share of dividends in total income is only 3.9%, a number that clearly understates the share of income from tradable assets. In order to account for this additional asset income, Heaton and Lucas adjust the parameters of the Markov chain so that the average share of dividends is 15%; specifically, they multiply the dividend share by a constant.

⁷Letting λ denote the growth rate of total income and η denote the share of dividends in total income, this process implies that that $E(\lambda) = 1.019$, $E(\eta) = 0.15$, $\text{s.d.}(\lambda) = 0.028$, $\text{s.d.}(\eta) = 0.008$, $\text{corr}(\lambda, \eta) = 0.29$, $\text{corr}(\lambda, \lambda_{-1}) = 0.18$, and $\text{corr}(\eta, \eta_{-1}) = 0.69$.

match the observed risk-free rate and equity premium. The Gini coefficient for wealth in our model turns out to equal $1 - \theta$. The Gini coefficient for total income in state i is $\psi_i(1 - \eta_i) + \eta_i - \theta$. Ignoring the (negligible) nonlinearity in this equation, the average Gini coefficient for income is $\psi(1 - \eta) + \eta - \theta$, where ψ and η are, respectively, the average share of labor income that goes to the big guys and the average dividend/output share. As reported in Díaz-Giménez, Quadrini, and Ríos-Rull (1997), the wealth Gini in U.S. data is 0.78 and the income Gini in U.S. data is 0.57, which implies that we should set θ equal to 0.22 and ψ to 0.75. In the experiments reported below, whenever we vary θ in order to show the effects of changes in wealth concentration, we simultaneously vary income concentration (by varying ψ) so that the income Gini is equal to a constant fraction of the wealth Gini, with this fraction being the one observed in the data, i.e., $0.57/0.78$.

Having pinned down the average value of the ψ_i 's (so as to match the income Gini), we need to select its stochastic properties. Recall that the purpose of this is to allow the big guy's income process to be more variable and correlated with dividends. This is a priori reasonable and will allow us to affect asset prices in the direction of delivering a higher equity premium. Ideally one would use time series properties of the income Gini to calibrate the properties of the ψ process, but we do not know of time series evidence for the Gini. Instead, we use a very simple, one-dimensional method: we tie the variations in ψ to θ . We assume, according to a specific formula, that a decrease in θ —more wealth concentration—will be associated with a ψ that has a higher covariance with η . Specifically, we assume that the growth in the big guy's endowment process is equal to aggregate income growth times $\frac{\eta' + \theta(1 - \eta')}{\eta + \theta(1 - \eta)}$. When $\theta = 1$, this expression is 1 and the big guy's endowment growth equals aggregate income growth, so it varies relatively little and covaries little with dividend growth. When θ is 0, on the other hand, his endowment growth equals aggregate income growth times growth in the labor share. The labor share covaries positively with dividend growth, so this delivers what we seek: more variation, and a higher covariation with the asset returns, of the big guy's growth process.

We should also show, for comparison, what the data says about our variables of interest. Our source here is Cecchetti, Lam, and Mark (2000). Let r_e be the annual (net) return on equity (expressed as a percentage), let r_f be the annual (net) return on the risk-free asset (expressed as a percentage), and let $r_d = r_e - r_f$ be the difference between these two returns. Let μ_i , $i = e, f, d$, and σ_i , $i = e, f, d$, denote the means and standard deviations, respectively,

of these three variables. Using annual data for 1889–1985, Cechetti, Lam, and Mark report the following estimates: $\mu_d = 5.75$, $\mu_f = 2.66$, $\sigma_d = 19.02$, $\sigma_f = 5.13$, and $\text{corr}(r_d, r_f) = -0.24$. From these estimates, one can deduce that $\mu_e = 8.42$, $\sigma_e = 18.47$, and $\sigma_f = 5.13$. Moreover, the market price of risk (μ_d/σ_d) is 0.30.

4.2 Findings

4.2.1 The effects of risk aversion with one vs. with two agents

We first set $\delta = 0.99$, $\theta = 1$, and $\beta = 1$ to obtain a set of results regarding the effect of risk aversion in the “standard model”:

Table 1

α	μ_e	μ_f	μ_d	σ_e	σ_f	μ_d/σ_d
0	2.90	2.82	0.076	2.87	0.59	0.027
-2	6.48	6.32	0.156	2.81	1.81	0.073
-4	9.77	9.63	0.138	3.81	3.11	0.063
-6	12.75	12.73	0.025	5.29	4.45	0.009
-8	15.41	15.59	-0.180	6.92	5.82	-0.048
-10	17.74	18.21	-0.469	8.61	7.21	-0.098

Note here that the equity premium increases and then *decreases* as risk aversion increases, eventually becoming negative for large enough values of risk aversion. This is a result of the properties of the particular driving process.⁸

Moving to our two-agent model, we notice that the nonmonotonic relationship between the equity premium and risk aversion goes away when we lower θ (the fraction of big guys). In particular, when we set $\theta = 0.3$, we obtain the results in Table 2. (For the experiments in Table 2, the wealth Gini equals 0.7, or roughly 90% of the observed wealth Gini. As discussed in Section 4.1, we therefore set the income Gini equal to 90% of its observed value, or 0.51. In this case, the ψ_i 's increase from 0.77 to 0.79 across the four states, with ψ , the average of the ψ_i 's, being 0.78.)

⁸The wealth Gini equals 0 when θ equals 1, so, as discussed in Section 4, we also set the income Gini to 0 by setting each of the ψ_i 's equal to 1.

Table 2

α	μ_e	μ_f	μ_d	σ_e	σ_f	μ_d/σ_d
0	2.91	2.81	0.105	3.31	0.17	0.032
-2	6.49	6.18	0.312	3.40	0.51	0.093
-4	9.73	9.22	0.510	3.61	0.87	0.146
-6	12.60	11.90	0.699	3.92	1.24	0.188
-8	15.08	14.20	0.882	4.31	1.60	0.221
-10	17.18	16.11	1.066	4.77	1.97	0.246

These results are dramatically different. First, decreasing θ increases the risk premium significantly for a given level of risk aversion: the return on equity does not change much, but the return on bonds goes down. Secondly, the variability of the bond return falls quite dramatically. Third, and in conclusion, we see that the effect of increased risk aversion on the price of risk is now monotone (increasing) and strong: it is thus the labor income process and low variability that makes the risk premium decreasing.

4.2.2 The effects of wealth concentration

Motivated by the findings above, let us display the effects of changing θ when $\delta = 0.99$, $\alpha = -2$, and $\beta = 1$ (i.e., still the standard model except for the heterogeneity):

Table 3

θ	μ_e	μ_f	μ_d	σ_e	σ_f	μ_d/σ_d	G_w	G_i
1	6.48	6.32	0.156	2.81	1.81	0.073	0	0
0.3	6.49	6.18	0.312	3.40	0.51	0.093	0.70	0.51
0.1	6.64	5.90	0.739	6.77	1.94	0.114	0.90	0.66
0.01	7.00	5.44	1.556	11.32	4.18	0.148	0.99	0.72

(In the tables, G_w is the wealth Gini and G_i is the income Gini.) We see that $\theta = 0.01$ delivers a risk premium of about one-and-a-half percent and a market price of risk of about 0.15.⁹

⁹In Tables 3 and 4, the ψ_i 's range from 0.64 to 0.74 (with an average of 0.69) when $\theta = 0.01$; they range from 0.69 to 0.75 (with an average of 0.71) when $\theta = 0.1$; and they range from 0.77 to 0.79 (with an average of 0.78) when $\theta = 0.3$. When $\theta = 1$, the ψ_i 's are all equal to one.

We now set $\alpha = -4$ (keeping $\delta = 0.99$, and $\beta = 1$) and display the effects of varying wealth concentration:

Table 4

θ	μ_e	μ_f	μ_d	σ_e	σ_f	μ_d/σ_d	G_w	G_i
1	9.77	9.63	0.138	3.81	3.11	0.063	0	0
0.3	9.73	9.22	0.510	3.61	0.87	0.146	0.70	0.51
0.1	10.01	8.43	1.585	9.47	3.35	0.179	0.90	0.66
0.01	10.81	7.14	3.674	17.27	7.21	0.235	0.99	0.72

With increased risk aversion, we see that $\theta = 0.01$ delivers a risk premium of nearly 4% and a market price of risk of about 0.24!

4.2.3 The urge to save

We now set $\delta = 0.99$, $\theta = 1$, $\gamma = 3$, and $\alpha = -4$ and vary β (recall that a higher β increases the urge to save):

Table 5

β	μ_e	μ_f	μ_d	σ_e	σ_f	μ_d/σ_d
0.5	13.92	13.76	0.157	3.84	3.29	0.080
1.0	9.77	9.63	0.138	3.81	3.11	0.063
2.0	6.30	6.18	0.123	3.86	2.95	0.050
4.0	4.20	4.08	0.115	3.92	2.86	0.043
6.0	3.45	3.34	0.112	3.95	2.82	0.041
8.0	3.07	2.96	0.111	3.97	2.81	0.040
10.0	2.84	2.73	0.110	3.98	2.80	0.039

Naturally, an increase in β decreases the returns. In particular, in this table, we are keeping the degree of risk aversion constant, and we see a sharp fall in the return on stock as well. The equity premium and the market price of risk both fall as well. This effect is very similar to that obtained from a decrease in the discount rate in the standard model. See Kocherlakota (1990) for details.

4.2.4 Matching the moments

The preceding sections show how risk aversion (α), the “urge to save” (β), and wealth concentration (θ) affect the first and second moments of equilibrium asset prices. In this section, we fix θ at three different values (0.3, 0.22, 0.1, and 0.01) and choose α and β so as to match (approximately) the risk-free rate and the equity premium observed in U.S. data.¹⁰ As discussed in Section 4.1, the Gini coefficient of the wealth distribution in the model is $1 - \theta$, which suggests that $\theta = 0.22$ is a reasonable choice for θ (since the Gini coefficient of the U.S. wealth distribution is 0.78).¹¹ The results are as follows ($\delta = 0.99$ and $\gamma = 3$ in this table):

Table 6

θ	α	β	μ_e	μ_f	μ_d	σ_e	σ_f	μ_d/σ_d	G_w	G_i
0.30	-33.0	3.0	8.53	2.79	5.75	14.39	6.15	0.44	0.70	0.51
0.22	-21.4	3.8	8.50	2.71	5.79	14.08	5.88	0.45	0.78	0.57
0.10	-10.5	3.3	8.40	2.51	5.89	17.42	7.18	0.37	0.90	0.66
0.01	-5.4	2.3	8.42	2.80	5.62	20.94	8.58	0.30	0.99	0.72
	Data		8.42	2.66	5.75	18.47	5.13	0.30	0.78	0.57

In addition to matching (by design) the first moments of the asset prices, the four parameterizations of the model economy reported in this table do a reasonable job of matching the three (unconditional) second moments of the asset prices.

5 Conclusions and final remarks

We have solved a Gul-Pesendorfer asset-pricing model where the temptation is one of savings urges among some investors. This model allows us to depict the world of consumers/investors in a not altogether unrealistic fashion: there is a small group of investors who are very rich and a large group of consumers who are poor and, as we model it here in order to simplify, borrowing-constrained. The rich investors are rich precisely because of their

¹⁰It seems clear that, for each of the four values of θ that we consider, we would be able to match *exactly* the risk-free rate and the equity premium by varying α and β appropriately.

¹¹When $\theta = 0.22$, the ψ_i 's vary from 0.74 to 0.77, with an average value of 0.75.

attitudes toward savings: with a single-minded focus on wealth accumulation, they end up dominating the economy in asset holdings. As such, they bear more risk, and their intertemporal preferences will be those determining asset prices, and not those of the “average consumer”. The average consumer, in contrast, will feature “hand-to-mouth” behavior. One implication that our setup has is that the risk-free rate becomes naturally lower. Secondly, with a reasonably high degree of risk aversion and concentrated risk, the equity premium can also be made much larger than in the standard model. Thirdly, we show that our parameterization also delivers low volatility of the risk-free rate.

Is our preference parameterization a sensible one? Clearly, the idea that investors have an urge to save is a very controversial one. But how strong an urge do we need in our “successful” parameterizations, i.e., in those that do a decent job of matching the asset price moments? We investigated the following comparison: suppose the agent (the stock investor, in this case) could choose between operating in the equilibrium we calculated, where he is subject to temptation and self-control costs, and one where he had no choices, but simply consumed the equilibrium allocation scaled down by a common factor at all dates and states. This factor would thus have to be less than one, because the equilibrium we compute involves self-control costs. What value of this factor would make the investor indifferent between the two environments? It turns out that the value is quite close to 1 for the parameterization above that gives the best match to the data: it is 0.997, thus indicating a self-control cost of less than a half a percent of consumption. With lower risk aversion, the cost goes up, however, and with logarithmic utility (keeping the other parameter values fixed) the cost is 7%. We also know that if γ , the parameter regulating the strength of temptation, were taken to infinity, the model would reduce to the Laibson model. For this limiting case, we proved in Krusell, Kuruşçu, and Smith (2001) that the disutility of self-control has to go to zero (independently of the value of risk aversion).

Another way of evaluating the importance of the urges to save can be given by comparing equilibrium utility to that given by another no-choice allocation, but not the equilibrium allocation scaled down as in the previous comparison. Instead, consider the best feasible allocation for the consumer—in the sense of being inside the budget—from the point of view of commitment utility (as given by \tilde{U}). This allocation is also available in equilibrium, but because of temptation it is not selected. However, if it were forced upon

the agent, it would make him better off by definition. Our measure here of how costly the urges are is thus by how much this alternative, no-choice path would have to be scaled down in order to deliver equilibrium utility. We find, again for our best parameterization, that the cost is now up to about 4%.¹² This measure, finally, is increasing in γ : as the temptation grows stronger, the disutility of self-control goes to zero, but the overall cost of urges of course goes up.

A comparison with Kocherlakota’s 1990 paper is also useful: is our framework with self-control costs from savings urges more reasonable than the one he proposes, where there are no self-control costs but the discount rate is above one? One could consider how the agents of the two kinds would compare consumption today and consumption in the distant future under various assumptions on the growth rate of consumption. If the growth rate of consumption is two percent, the intertemporal elasticity of substitution is $1/2$, and the discount rate in a Kocherlakota economy is such that an agent would choose constant consumption at a 1% rate of interest, then his agent would be willing to pay 0.37 units of consumption today for 1 additional unit in 100 years—this is 1.01^{-100} , as it should be. Our agent, in contrast, would only be willing to pay 0.007 units today, under our kind of calibration (with a “short-run” discount rate larger than one and a “long-run” rate below one) and assuming that the the comparison is between the two “no-choice” paths $\{c, 1.02c, 1.02^2c, \dots, 1.02^{100}c + 1, 1.02^{101}, \dots\}$ and $\{c + x, 1.02c, 1.02^2c, \dots, 1.02^{100}c, 1.02^{101}, \dots\}$, where x thus defines the current willingness to pay for 1 unit in 100 years. That is, our agent would be much more impatient in this comparison, but perhaps not unreasonably so. In contrast, suppose we considered, for the same preference parameter values, a zero consumption growth rate. Then our agent would pay 0.37 units today for one unit in 100 years, whereas the Kocherlakota agent would now be willing to pay 19.4 units! Thus, depending on the context, the two setups look more or less reasonable; perhaps ours is more robust in the sense of never producing crazy numbers.

Obviously, as a general point one needs to learn more about the actual behavior of the large investors. We know of no empirical study that is helpful in this regard. Our approach here can perhaps best be characterized as overly bold theorizing, but we do believe that moving in the direction of

¹²Here as well, the cost is decreasing in risk aversion; for logarithmic utility, it would be 41%.

understanding the “psychology” of large investors is very important both for understanding asset pricing and macroeconomics.

REFERENCES

Basak, Suleyman and Domenico Cuoco (1998), “An Equilibrium Model with Restricted Stock Market Participation,” *Review of Financial Studies* 11, 309–341.

Cecchetti, Stephen G., Pok-Sang Lam, and Nelson C. Mark (2000), “Asset Pricing with Distorted Beliefs: Are Equity Returns Too Good to be True?”, *American Economic Review* 90, 787–805.

Dai, Qiang (2001), “Asset Pricing in a Neoclassical Model with Limited Participation,” manuscript (New York University).

Díaz-Giménez, Javier, Vincenzo Quadrini, and José-Víctor Ríos-Rull (1997), “Dimensions of Inequality: Facts on the U.S. Distribution of Earnings, Income, and Wealth,” *Federal Reserve Bank of Minneapolis Quarterly Review* 21, 3–21.

Epstein, Larry G. and Stanley E. Zin (1989), “Substitution, Risk Aversion, and the Temporal Behavior of Consumption Growth and Asset Returns: A Theoretical Framework,” *Econometrica* 57, 937–969.

Gul, Faruk and Wolfgang Pesendorfer (2000a), “Temptation and Self-Control,” forthcoming in *Econometrica*.

Gul, Faruk and Wolfgang Pesendorfer (2000b), “Self-Control and the Theory of Consumption,” manuscript (Princeton University).

Guvenen, Muhammet F. (2001), “Mismeasurement of the Elasticity of Intertemporal Substitution: The Role of Limited Stock Market Participation,” manuscript (Carnegie Mellon University).

Heaton, John and Deborah J. Lucas (1996), “Evaluating the Effects of Incomplete Markets on Risk sharing and Asset Prices,” *Journal of Political Economy* 104, 443–487.

Kocherlakota, Narayana (1990), “On the Discount Factor in Growth Economies,” *Journal of Monetary Economics* 25, 43–47.

Krusell, Per, Burhanettin Kuruşçu, and Anthony A. Smith, Jr. (2001),

“Temptation and Taxation,” manuscript (University of Rochester and Carnegie Mellon University).

Krusell, Per, Burhanettin Kuruşçu, and Anthony A. Smith, Jr. (2000), “Asset Pricing with Laibson-Epstein-Zin Preferences,” manuscript (University of Rochester and Carnegie Mellon University).

Krusell, Per and Anthony A. Smith, Jr. (1997), “Income and Wealth Heterogeneity, Portfolio Choice, and Equilibrium Asset Returns,” *Macroeconomic Dynamics* 1, 387–422.

Krusell, Per and Anthony A. Smith, Jr. (1998), “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy* 106, 867–896.

Laibson, David (1997), “Golden Eggs and Hyperbolic Discounting,” *Quarterly Journal of Economics* 112, pp. 443-477.

Lucas, Robert E., Jr. (1978), “Asset Prices in an Exchange Economy,” *Econometrica* 46, 1429–1445.

Luttmer, Erzo G.J. and Thomas, Mariotti (2000), “Subjective Discounting in an Exchange Economy,” manuscript (London School of Economics).

Mehra, Rajnish and Edward C. Prescott (1985), “The Equity Premium: A Puzzle,” *Journal of Monetary Economics* 15, 145–161.

Tauchen, George and Robert Hussey (1991), “Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models,” *Econometrica* 59, 371–396.

Appendix

This Appendix explains how we compute the 32 coefficients in the actual decision rules, given a set of 16 asset prices. This Appendix also explains how we determine the 16 market-clearing prices.

The basic idea of the algorithm for computing the decision rules coefficients is to take as given an initial set of coefficients, use equation (13) to update this set of coefficients, and then continue iterating until convergence. Let a'_{ij} and b'_{ij} , $i = 1, \dots, 4$ and $j = 1, \dots, 4$, denote the initial set of coefficients; these coefficients are used to determine tomorrow's (actual) asset decisions (hence the "primes" on the coefficients). We now describe how we use equation (13) to compute the coefficients that determine today's (actual) asset decisions, taking as given the a'_{ij} 's and b'_{ij} 's.

First, note that equation (12) can be used to solve for the ratio of temptation consumption to actual consumption in the next period, assuming that next period's (actual) decisions are determined by the a'_{ij} 's and b'_{ij} 's. Because this is the only way in which temptation decisions enter equation (13), we do not need to consider separately the first-order condition for temptation behavior. In particular, it is straightforward to show that

$$\frac{\tilde{c}'_j}{c'_j} = \left(1 - \left(\sum_{l=1}^4 q_{jl} b'_{jl} \right) \left(1 - \left(\frac{1 + \beta\gamma}{(1 + \gamma)\beta} \right)^{\frac{1}{\alpha-1}} \right)^{-1} \right).$$

Next, define

$$K_{ij} = \left(\frac{\delta(1 + \beta\gamma)}{1 + \gamma} \right)^{\frac{1}{\alpha-1}} \left(\frac{\pi_{ij}}{q_{ij}} \right)^{\frac{1}{\alpha-1}} g_{ij} \left(1 + \gamma - \gamma \left(\frac{\tilde{c}'_j}{c'_j} \right)^{\alpha-1} \right)^{\frac{1}{\alpha-1}}.$$

K_{ij} will be treated as a constant (which depends on the b'_{ij} 's) when solving for today's (actual) decision rule coefficients.

The Euler equation (13) can now be written:

$$c - K_{ij} c'_j = 0,$$

where

$$\begin{aligned} c &= \hat{\omega} - \sum_{l=1}^4 q_{il} \hat{s}'_{il}(\hat{\omega}) \\ &= - \sum_{l=1}^4 q_{il} a_{il} + \left(1 - \sum_{l=1}^4 q_{il} b_{il} \right) \hat{\omega} \end{aligned}$$

and

$$\begin{aligned} c'_j &= \hat{\omega}'_j - \sum_{l=1}^4 q_{jl} \hat{s}'_{jl}(\hat{\omega}'_j) \\ &= a_{ij}/g_{ij} + (b_{ij}/g_{ij}) \hat{\omega} + 1 - \sum_{l=1}^4 q_{jl} \left(a'_{jl} + b'_{jl} (a_{ij}/g_{ij} + (b_{ij}/g_{ij}) \hat{\omega} + 1) \right) \end{aligned}$$

Inserting these expressions into the Euler equation and combining terms yields:

$$\begin{aligned} -K_{ij} \left(1 - \sum_{l=1}^4 q_{jl} (a'_{jl} + b'_{jl}) \right) - \sum_{l=1}^4 q_{il} a_{il} - (K_{ij}/g_{ij}) \left(1 - \sum_{l=1}^4 q_{jl} b'_{jl} \right) a_{ij} + \\ \left(1 - \sum_{l=1}^4 q_{il} b_{il} - (K_{ij}/g_{ij}) \left(1 - \sum_{l=1}^4 q_{jl} b'_{jl} \right) b_{ij} \right) \hat{\omega} = 0. \end{aligned}$$

In other words, the Euler equation is an affine function of $\hat{\omega}$. Since the Euler equation must hold for all values of $\hat{\omega}$, both the intercept and the slope in this equation must equal zero. These conditions determine the a_{ij} 's and the b_{ij} 's. In particular, for each $i = 1, \dots, 4$, the set of coefficients $\{a_{ij}\}_{j=1}^4$ must solve the following linear system of equations for $j = 1, \dots, 4$:

$$-K_{ij} \left(1 - \sum_{l=1}^4 q_{jl} (a'_{jl} + b'_{jl}) \right) - \sum_{l=1}^4 q_{il} a_{il} - (K_{ij}/g_{ij}) \left(1 - \sum_{l=1}^4 q_{jl} b'_{jl} \right) a_{ij} = 0.$$

In addition, for each $i = 1, \dots, 4$, the set of coefficients $\{b_{ij}\}_{j=1}^4$ must solve the following linear system of equations for $j = 1, \dots, 4$:

$$1 - \sum_{l=1}^4 q_{il} b_{il} - (K_{ij}/g_{ij}) \left(1 - \sum_{l=1}^4 q_{jl} b'_{jl} \right) b_{ij} = 0.$$

The 32 actual decision rule coefficients, given the a'_{ij} 's and b'_{ij} 's, can therefore be computed by solving 8 separate sets of linear equations, each of which involves 4 of the coefficients.

Having computed the a_{ij} 's and b_{ij} 's that set the Euler equation to zero for all values of $\hat{\omega}$, we then check whether these coefficients are sufficiently close to the ones taken as given (i.e., the a'_{ij} 's and b'_{ij} 's). If not, then we replace the a'_{ij} 's and b'_{ij} 's with the new coefficients and continue iterating until our convergence criterion is satisfied. In particular, we require that

$$\max \left(\max_{i,j} |a_{ij} - a'_{ij}|, \max_{i,j} |b_{ij} - b'_{ij}| \right)$$

be less than 10^{-8} . This algorithm is fairly robust to different choices for the initial a'_{ij} 's and b'_{ij} 's, including zero. In the interest of speed, however, it is generally a good idea to use an initial guess that comes from a parameterization of the model that is “close” to the one under consideration.

To find the market-clearing prices, we use Newton’s algorithm. We use numerical derivatives to compute the Jacobian of the mapping from the q_{ij} ’s to excess demands in each of the markets (i.e., $a_{ij} - b_{ij}$ for each i and j). During the first few iterations of Newton’s algorithm, we find that it is necessary (unless the initial guess for the prices happens to be very close to the market-clearing prices) to take smaller steps than those dictated by Newton’s algorithm. We do so by letting the prices for the next iteration be a weighted average of the old prices and the new prices that would be selected if we used a full step. When $\beta = 1$, it is straightforward to show that the equilibrium prices $q_{ij} = \delta\pi_{ij}g_{ij}^{\alpha-1}$. These prices serve as a good initial guess for the prices that obtain when we perturb β away from 1. The converged prices for such a perturbation serve, in turn, as a good initial guess for the next perturbation of β .