

ARTICLES

THE WEALTH DISTRIBUTION AND THE DEMAND FOR STATUS

YULEI LUO

University of Hong Kong

ERIC R. YOUNG

University of Virginia

Standard economic theories of asset markets assume that assets are valued entirely for the consumption streams they can finance. This paper examines the introduction of the demand for status (as a function of wealth) into a model of uninsurable idiosyncratic risk—the “spirit of capitalism” (“soc”) assumption. We find that soc preferences lead to less inequality in wealth; placing wealth into the utility function leads to a shrinking wealth distribution. The drop in wealth concentration is smaller if the utility function implies status is a luxury good, but no parametrization leads to higher wealth Gini coefficients than the benchmark case. We then consider the consequences of revenue-neutral tax reforms with and without soc preferences, finding that they make little difference for this policy experiment.

Keywords: Wealth Distribution, Spirit of Capitalism, Preference for Status

1. INTRODUCTION

Our interest in this paper is to evaluate the implications for the wealth distribution of a particular modification to preferences: we assume agents value wealth directly. We consider a simple specification in which agents value wealth because it gives them status; this specification has come to be known as the “spirit of capitalism” or “soc” assumption and dates at least to Duesenberry (1949). As in Bakshi and Chen (1996), we let status be a function of individual wealth, giving agents an additional motivation to accumulate assets. The standard specification of preferences—a model based on Aiyagari (1994) with elastic labor supply—generates saving only for precautionary purposes; because the equilibrium interest rate must be strictly below the rate of time preference, agents slow down their accumulation of wealth as soon as they become sufficiently well insured. Our model adds another effect—higher wealth confers utility directly as well as indirectly through the consumption purchases that wealth can finance.¹ As we will show, this motive can generate wealth accumulation even in the presence of negative real returns.

The authors would like to thank the helpful comments of Wouter den Haan, Wojciech Kopczuk, Chris Sims, the editor William Barnett, and an anonymous referee. Any errors remaining are, of course, the responsibility of the authors. Address correspondence to: Dr. Eric R. Young, Department of Economics, University of Virginia, Charlottesville, VA 22904, USA; e-mail: ey2d@virginia.edu.

Spirit of capitalism preferences have been shown to improve asset-pricing models by Bakshi and Chen (1996), Smith (2001), and Kenc and Dibooglu (2007), to affect long-run growth by Zou (1994), to change the implications of taxation for growth by Gong and Zou (2002), and to change the model predictions for international risk sharing by Chue (2004).² In work more closely related to this paper, Carroll (2002) examines the implications of valuing wealth directly—albeit only after one dies, leaving a bequest—for the portfolio decisions of households, finding that this specification is better able to match the portfolios of the rich; Reiter (2004) studies the related issue of the overall savings rate of the wealthiest individuals, again finding some improvement from introducing soc preferences. The most closely related work is De Nardi (2004), which studies the effects of valuing bequests for wealth concentration. We will comment more on this paper after we describe our results, because they will be quite different.

Our paper is ultimately motivated by policy considerations, although our focus will not be primarily on policy applications. A huge literature studies policy questions in the Aiyagari (1994) model, focusing on a wide range of possible issues. We want to draw particular attention to two papers that do not directly conduct policy experiments but have policy relevance: Krusell and Smith (2002) and Dávila, Hong, Krusell, and Ríos-Rull (2006). The first paper asks how the distribution of wealth would change if the business cycle were eliminated (without explicitly modelling how that elimination is achieved). The effect on the concentration of wealth depends on how the model achieves a good fit to the U.S. Gini coefficient. If one uses heterogeneity in discount factors to produce high wealth concentration—the agents who draw high values will save a lot while those who draw low values will consume—eliminating the business cycle will generate more wealth concentration; by weakening the precautionary savings motive the elimination of cycles reduces the savings of the low discount factor types, leading to higher returns and therefore more wealth accumulation by the high types. An alternative approach to match concentration is to assume that earnings is mismeasured—the most commonly used data for earnings shocks, the PSID, simply does not contain enough of the wealthy agents to accurately measure the process for earnings.³ Under this approach eliminating cycles will result in a smaller Gini coefficient for wealth, not a larger one [at least, that outcome is the conjecture of Krusell and Smith (2002)].

The second paper asks how a social planner who is constrained to respect individual budget constraints would alter the distribution of wealth; since the poor have the highest marginal utility of wealth the planner's choices will be geared towards increasing their consumption. If the idiosyncratic risk faced by agents is driven primarily by changes in their wages, the asset-poor will have more labor income than capital income; in this case the planner would want to increase the capital stock to drive up wages and provide additional insurance. If instead risk is modeled as unemployment shocks, the asset-poor will have essentially no labor income, meaning their income is composed entirely of their (low) capital income. As a result the planner would want less capital in order to raise the

return and make capital a more effective insurance vehicle.⁴ Again, the specific assumptions of the model have qualitatively-important consequences; since the model with unemployment shocks can match the concentration of wealth only when combined with discount factor heterogeneity but the model with earnings shocks can potentially do so without these shocks, the optimal amount of aggregate capital will differ across specifications that can match the U.S. distribution. These differences may change the nature of optimal capital taxation, a point which we will discuss more later.⁵

Why do we explicitly discuss these two papers as motivation? Our *ex ante* belief was that spirit of capitalism preferences would be a force for wealth concentration, particularly if they imply that status is a luxury good, because they would tend to increase the desire of wealthy agents to get even wealthier; if this increased demand for assets were not shared equally across the population, as it would not be for nonhomothetic preferences, the equilibrium decline in returns would lead to less accumulation by the poor and therefore more concentration. The question at hand would then be whether matching the distribution of wealth using soc preferences yielded different implications for policy experiments, as Krusell and Smith (2002) and Dávila et al. (2006) found for their particular experiments. The question is of course predicated on being able to match the wealth distribution in a model with soc preferences, which is the main focus of this paper; we find that these preferences have little potential at resolving the excess concentration of wealth puzzle. This finding does not render moot the question of how alternative models of wealth concentration affect policy implications because we are able to find versions of the soc model that match the benchmark model's wealth distribution exactly.

We find that soc preferences are a force *against* wealth concentration, rather than for it; in fact, they tend to lead to a collapse in the wealth distribution. As we increase the coefficient that controls the desire for status, we must simultaneously decrease the discount factor to keep aggregate wealth from rising. Two effects thus contribute to a collapse of the wealth distribution—the minimum level increases due to the infinite marginal utility of wealth at zero while the maximum level falls due to the drop in the discount factor. As the strength of this effect goes to infinity, the wealth distribution collapses to a single point. If the discount factor is not adjusted, the distribution still collapses but does so around a much higher mean—the upper bound on wealth increases at a slower rate than the lower bound. The collapse of the distribution is mitigated if the utility function implies that status is a luxury good. We are unable to provide any parameter combinations which lead to absolute increases in the Gini coefficient of wealth over the benchmark; we are able to find combinations that imply essentially the same wealth distribution, namely those with highly nonhomothetic preferences. We conclude from these experiments that the spirit of capitalism is a force against wealth concentration.⁶

Our results are therefore in opposition to the findings in De Nardi (2004). Because her specification implies different preferences for different agents—only

retirees face mortality risk and value bequests—it can generate more wealth concentration. Our model produces the same result if we assume only some agents have soc preferences. The reason that a model with heterogeneous soc preferences (or bequests valued only by retirees) can produce wealth concentration is that it performs very similarly to a model with heterogeneous discount factors, as in Krusell and Smith (1998); the high discount factor types correspond to those agents with increased demand for savings (the soc preference type). The presence of these households decreases the equilibrium return to savings and thus reduces the savings of other types of households, leading to increased wealth concentration.

In the final section of the paper, we explore the implications of soc preferences for tax reform. Given that the preferences we consider are capable of producing the same Gini coefficients on wealth as standard ones, it is important to assess whether the predictions for the effects of policies are also similar, in light of the results in Krusell and Smith (2002) and Dávila et al. (2006). We calibrate the model to the U.S. progressive tax system and then compute two reforms, replacing the progressive tax with a flat income tax and with a flat consumption tax (both experiments are revenue-neutral). We compute the welfare consequences of these reforms (including the costs of the transitional dynamics), finding that on average households would prefer the consumption tax to the income tax reform independent of their preferences; this result has a long history in public finance although it is new in the context of this model.⁷ The individual welfare gains from the income tax reform are increasing in initial capital and increasing in initial productivity, whereas for the consumption tax reform they are decreasing in initial capital and increasing in initial productivity. Interestingly, we find that income tax reforms are Pareto-improving but consumption tax reforms are not, despite the fact that the consumption tax reform dominates on average. As the strength of the demand for status increases, welfare gains initially fall but then rise for both types of tax reform, but the changes are small; thus, soc preferences do affect the welfare calculations, but not in a dramatic way.

The result that flat income taxation Pareto-dominates progressive taxation but flat consumption taxation does not seem to be new. The operative difference (in a model with standard preferences) between the income tax and the consumption tax is the taxation of capital; constant consumption taxes distort the same margin as labor income taxes, namely, the equality between the wage and the marginal rate of substitution between consumption and leisure, but leave the intertemporal margin undistorted. Thus, our results echo the results in Aiyagari (1995) and Conesa, Kitao, and Krueger (2007) on the optimality of positive taxes on capital when asset markets are incomplete (because the economy tends to save too much, capital taxation can reduce capital to an appropriate level); however, the fact that consumption taxes are better on average seems to be a new result for this class of models, as is the result that they harm the initially wealthy. With soc preferences, this decomposition is more difficult because consumption taxes also distort an intertemporal margin, similar to a tax on capital income, because they raise the price of status. Nevertheless, for modest departures from standard preferences

that imply the same initial wealth distribution we find that soc preferences do not change the welfare rankings. Interestingly, as we raise the parameter that governs the strength of the soc preference, we find that costs change in a nonmonotonic manner; initially, the welfare gains from tax reform are decreasing in this parameter but eventually they increase, becoming larger than the benchmark. The quantitative differences between the parameter settings are small, however, so we conclude that tax reform experiments are unlikely to be sensitive to the soc assumption.

2. THE MODEL ECONOMY

Our model economy will feature partially uninsurable idiosyncratic wage risk and markets will be exogenously incomplete; we will allow households to hold only aggregate capital for savings, and holdings of this asset are restricted by an exogenous borrowing limit. However, the household supplies labor elastically which aids in smoothing consumption.

2.1. The Environment

We consider a model economy with a large (measure 1) population of infinitely lived consumers, as in Aiyagari (1994).⁸ There is only one consumption good per period and we assume that all agents have the same preferences over streams of consumption, social status, and leisure, represented by

$$E_0 \sum_{t=0}^{\infty} \beta^t u(c_t, s_t, l_t), \quad (2.1)$$

where $u_s(c_t, s_t, l_t) > 0$ (higher status is strictly preferred), $u_{ss}(c_t, s_t, l_t) < 0$ as discussed in Robson (1992), and the cross-derivative is unrestricted. We will confine ourselves to the class of utility functions which satisfy constant relative risk aversion over static gambles:

$$u(c_t, s_t, l_t) = \log(c_t) + \mu \log(l_t) + \frac{\theta}{1-\sigma} (s_t + \gamma)^{1-\sigma}; \quad (2.2)$$

our choice is dictated by the observation that the return to capital has been stationary over the postwar period. The status function $s_t = F(W_t, \bar{W}_t)$ is assumed to possess the following properties: (1) it is strictly increasing in the investor's absolute wealth at time t (W_t), so that higher wealth means higher status regardless of the wealth distribution for the group of people with whom the investor has social or professional contacts; (2) it is a decreasing function of the wealth of the social group to which this individual belongs (\bar{W}_t), implying $F_W > 0$ and $F_{\bar{W}} < 0$. In our economy, we will assume that W_t is the individual capital stock k_t ; because we focus mainly on stationary outcomes, we ignore the dependence on average wealth. In this utility function, $\sigma \geq 0$, $\theta \geq 0$, and $\gamma \geq 0$.⁹

Each agent is endowed with one unit of time with a stochastic productivity y_t . The budget constraint for the household is

$$c_t + k_{t+1} \leq (r_t + 1 - \delta)k_t + y_t w_t h_t, \quad (2.3)$$

where r_t is the rental rate on capital, $\delta \in [0, 1]$ is the depreciation rate, w_t is the wage rate, and h_t is labor supply. Because we intend to discuss the implications of tax policy where labor market distortions become critical, we assume agents value leisure in the initial model to ensure that our results do not only apply to a model with inelastic labor supply. Pijoan-Mas (2006) demonstrates that elastic labor supply has consequences for the degree of precautionary savings as well.

We assume that y_t is generated by a Markov process with stationary transitions described by a vector of realizations $\{y\}$ and transition probabilities $[\pi_{ij}]$. The time allocation constraint is

$$1 = h_t + l_t. \quad (2.4)$$

Capital is restricted to be nonnegative:

$$k_{t+1} \geq 0, \quad (2.5)$$

and claims contingent on the outcome of the idiosyncratic shock y are not marketed.¹⁰

The technology produces output Y_t as a Cobb-Douglas function of capital input K and labor input N :

$$Y_t = K_t^\alpha N_t^{1-\alpha}.$$

Output can be transformed into future capital K_{t+1} and current consumption C_t according to

$$C_t + K_{t+1} = Y_t + (1 - \delta)K_t. \quad (2.6)$$

2.2. The Market Arrangement

Consumers collect income from working and from renting the services of their capital. If the total amount of capital in the economy is denoted K and the total amount of labor supply is denoted H , the CRTS production function implies that the relevant first order conditions are

$$w(K, N) = (1 - \alpha) \left(\frac{K}{N} \right)^\alpha, \quad (2.7)$$

and

$$r(K, N) = \alpha \left(\frac{K}{N} \right)^{\alpha-1}. \quad (2.8)$$

We consider a recursive equilibrium definition, which includes a law of motion for the aggregate state of the economy as a key element. The aggregate state of the economy is the current measure (distribution) of consumers over holdings

of capital and productivity, which we denote by Γ . For the individual agent, the optimization problem can therefore be expressed as follows:

$$v(k, y) = \max_{1 \geq h \geq 0, c \geq 0, k' \geq k_b} \{u(c, s, 1 - h) + \beta E[v(k', y')|y]\}, \quad (2.9)$$

subject to

$$c + k' = r(K, N)k + w(K, N)hy + (1 - \delta)k \quad (2.10)$$

$$s = F(k), \quad (2.11)$$

and the stochastic law of motion for y . The decision rule for the updating of capital coming out of the problem is denoted by the function $\pi_k(k, y)$ and the one for labor is denoted $\pi_h(k, y)$.

DEFINITION 1. *A recursive competitive equilibrium is a value function $v(k, y)$, decision rules $\pi_k(k, y)$ and $\pi_h(k, y)$, pricing functions $r(\Gamma)$ and $w(\Gamma)$, and a law of motion $G(\Gamma)$ such that*

- (i) (v, π_k, π_h) solves the consumers' problems given prices and the law of motion;
- (ii) r and w are consistent with the firm's first-order conditions;
- (iii) G is generated by f , that is, by the appropriate summing up of agents' optimal choices for capital given their current state.
- (iv) The goods market clears: $C + K' = K^\alpha N^{1-\alpha} + (1 - \delta)K$.
- (v) Factor markets clear: $N = \int y \pi_h(k, y) \Gamma(k, y)$ and $K = \int k \Gamma(k, y)$

Our goal is generally to find stationary equilibria, so we need only the fixed point of the law of motion

$$\Gamma^* = G(\Gamma^*)$$

and need not compute the law of motion explicitly.¹¹ For transitional dynamics we compute $G(\Gamma)$ explicitly along a deterministic path.

3. RESULTS

We now present our results. Because the model will not produce a distribution with a known form [see Aiyagari (1994) or Young (2005) for discussions of the shape of the distributions produced], we use numerical methods to obtain our results; our algorithm is detailed in a computational appendix available upon request. This section contains four subsections. In the first subsection, we discuss the upper bound on asset returns implied by a model with soc preferences. We then provide results on increasing the strength of the demand for status in the homothetic case. The final two subsections present results for the two luxury good cases: $\gamma > 0$ and $\sigma < 1$.

3.1. Upper Bound for r

Our first result for the model is that the usual upper bound derived for the interest rate in an incomplete market model will hold here in a stronger form. With

complete markets, the steady state is defined by

$$r - \delta = \frac{1 - \beta \frac{u_s}{u_c}}{\beta} - 1;$$

the additional term $-\frac{u_s}{u_c}$ is the steady-state marginal rate of substitution between status and consumption. With the utility function we will choose (logarithmic preferences), the expression becomes

$$r - \delta = \frac{1 - \beta \theta \frac{c}{k+\gamma}}{\beta} - 1.$$

Because this additional term is positive, the steady-state interest rate is strictly lower than without status in the utility function; furthermore the long-run capital supply curve would be less than perfectly elastic even in complete markets. In our economy with incomplete markets, there is an upper bound on r implied by this equation; if we set $\theta = 0$ then we obtain the upper bound on r from Aiyagari (1994) that $r - \delta < \beta^{-1} - 1$. When $\theta > 0$ this bound does not hold, but the results in Carroll (2004) might be adapted to show that as k goes to infinity, $\frac{c}{k+\gamma}$ converges from above to some value $\chi < 1$, which implies that our bound would be

$$r - \delta < \frac{1 - \beta \theta \chi}{\beta} - 1 < \frac{1}{\beta} - 1. \quad (3.1)$$

χ must be less than one since $r - \delta$ is less than one and labor supply is zero for sufficiently high wealth.¹²

The intuition of the bound on r is straightforward. As noted in Aiyagari (1994), asset supply goes to ∞ as $r - \delta$ approaches $\beta^{-1} - 1$ from below, because any agent whose time rate of preference equals the market return would desire perfectly flat consumption; this desire requires an infinite amount of assets to satisfy in the presence of uninsurable risk. r must decline to keep asset supply finite, since capital demand is finite for any positive r . In our economy, asset supply will be finite for almost every agent, except possibly the one with the minimum $\frac{c}{k+\gamma}$, as this individual has the strongest desire to hold assets; therefore the bound on r keeps the asset supply of this agent finite.

3.2. Homothetic Preferences

In this subsection we discuss our baseline parametrization, where preferences are homothetic. We specify status as in Bakshi and Chen (1996):

$$s = k. \quad (3.2)$$

We set risk aversion over status to 1; the resulting preferences are represented by

$$u(c, s, l) = \log(c) + \mu \log(1 - h) + \theta \log(s).$$

For calibration, we choose β to match a capital/output ratio of 3.0, δ to match an investment/output ratio of 0.25, and μ to generate average hours of 0.3271 percent of the time endowment.¹³

Based on wage estimates in Floden and Lindé (2001), we specify $\log(y_i)$ to be

$$\log(y_i) = \omega_i + \epsilon_i \quad (3.3)$$

$$\omega'_i = \rho\omega_i + v'_i, \quad (3.4)$$

where $\epsilon_i \sim N(0, \sigma_\epsilon^2)$ is the transitory component and ω_i is the persistent component. The innovation term associated with ω_i is assumed to be distributed $N(0, \sigma_v^2)$. Floden and Lindé (2001) estimate $\rho = 0.913$, $\sigma_v^2 = 0.0426$, and $\sigma_\epsilon^2 = 0.0423$ for annual data. We then approximate this process with a seven-state Markov chain using the Tauchen (1986) procedure.

This process will not generate sufficient income inequality to replicate the observed Gini coefficient of wealth; however, it is sufficient to demonstrate how the parameter θ affects inequality. As a reference point, the data for the United States examined in Budría Rodríguez, Díaz-Giménez, Quadrini, and Ríos-Roll (2002) yields a Gini coefficient of wealth of 0.78 whereas our benchmark case ($\theta = 0$) generates a value of 0.55. Our model's prediction for the Gini coefficient is consistent with previous findings—such as Aiyagari (1994) and Floden and Lindé (2001)—that measured earnings cannot match the wealth distribution; the main discrepancy appears at the upper end of the wealth distribution where the model produces too few very wealthy individuals.¹⁴ The reason for this failure has been noted previously. Agents save only for precautionary purposes; since low assets imply that bad shocks will cause consumption instead of wealth to fall, households increase their savings to reduce the probability of reaching the liquidity constraint. When households become wealthy enough the probability of reaching the constraint is very small, meaning that these agents have no reason to continue saving; because the precautionary demand for savings declines rapidly with wealth, no very wealthy agents appear.

3.2.1. Wealth. Our concern in this subsection is the relationship between θ and the wealth distribution. Figure 1 presents the wealth distributions from several of the cases covered in Table 1.¹⁵ We can see from both the table and the figure that increasing θ has the following effects (keeping in mind that the wealth/income ratio and aggregate hours are kept constant):

1. The standard deviation of the wealth distribution shrinks;
2. The skewness of the wealth distribution converges to zero;
3. The kurtosis of the wealth distribution shrinks;
4. The Gini coefficient goes to zero.

Thus, we see that adding concern for wealth into the utility function of the household leads the wealth distribution to collapse. Figure 2 graphs the Lorenz curves for the model assuming various different values for θ . It is obvious that




FIGURE 1. Distribution of wealth.

when we increase the concern for the status the wealth distribution becomes more equal and converges to the 45° line. The contraction of the wealth distribution has several components. One, agents are risk averse toward gambles over status; that is, they prefer to smooth status, and by extension, wealth. Since the economy possesses an ergodic distribution and satisfies the mixing conditions discussed in Aiyagari and Alvarez (2001), each household understands that they will visit each point in the state space infinitely often over their infinite life; that is, the ergodic distribution is simultaneously the cross-sectional distribution at a point in time and the distribution of an individual over time. Risk aversion over wealth compels them to make this ergodic set smaller, which is exactly what they do. To get a handle on the size of this effect, we include in Table 1A the highest and lowest levels of wealth observed in our distribution (we carry the computation out to 10 decimal places, all of which we require to be zero) for each value of θ .¹⁶ The maximum value of k decreases rapidly as θ increases, lending support to the idea that households are reducing the size of their ergodic set. In addition, as θ increases the minimum value of k observed increases, as agents become less willing to hold low amounts of wealth. If we allow θ to go to infinity, the ergodic distribution converges to a single point.¹⁷

TABLE 1A. Moments of wealth distribution $\gamma = 0$ (wealth not a luxury good)

Models	Gini	Mean (k)	Std. (k)	Skew (k)	Kurt. (k)	Max (k)	Min (k)
$\theta = 0$	0.408	1.920	1.463	1.289	5.342	15.873	0.000
$\theta = 0.1$	0.163	1.910	0.554	0.473	3.315	5.653	0.242
$\theta = 0.5$	0.102	1.899	0.344	0.363	3.203	3.729	0.843
$\theta = 1.0$	0.086	1.894	0.288	0.323	3.147	3.248	0.963
$\theta = 5.0$	0.067	1.888	0.223	0.267	3.039	2.767	1.204

TABLE 1B. Moments of consumption distribution $\gamma = 0$ (wealth not a luxury good)

Models	Gini	Mean (c)	Std. (c)	Skew (c)	Kurt. (c)
$\theta = 0$	0.039	0.480	0.034	-0.386	5.068
$\theta = 0.1$	0.040	0.478	0.034	-0.040	3.080
$\theta = 0.5$	0.048	0.475	0.041	0.137	2.988
$\theta = 1.0$	0.052	0.474	0.043	0.191	2.981
$\theta = 5.0$	0.057	0.472	0.047	0.255	2.982

TABLE 1C. Moments of hours distribution $\gamma = 0$ (wealth not a luxury good)

Models	Gini	Mean (h)	Std. (h)	Skew (h)	Kurt. (h)	Corr (h, y)
$\theta = 0$	0.173	0.327	0.100	-0.529	3.072	0.866
$\theta = 0.1$	0.148	0.327	0.086	-0.458	3.107	0.872
$\theta = 0.5$	0.128	0.327	0.075	-0.490	3.359	0.833
$\theta = 1.0$	0.120	0.327	0.071	-0.509	3.498	0.812
$\theta = 5.0$	0.112	0.327	0.066	-0.553	3.782	0.770

The mean levels of wealth are not the same across the economies, despite the fact that the capital/output ratio and the total hours worked are forced to be equal, because there is a change in the distribution of hours. As noted in Pijoan-Mas (2006), shifts in the distribution of hours between low to high-productivity workers can alter the amount of precautionary savings. Table 1C presents the correlation between productivity and hours across values of θ . Starting from $\theta = 0$, the correlation initially rises with θ ; this increase continues until θ reaches a critical value somewhere between 0.1 and 0.5, where it begins to decline. When the demand for status is initially low, increasing it results in a shift of aggregate hours toward the productive, leading to higher aggregate productivity. As θ continues to rise, however, this relationship reverses itself, producing low productivity workers who supply a lot of labor to get their wealth and status to increase. However, when θ begins to get very large, the correlation again begins to rise; this occurs above $\theta = 5.0$.¹⁸




FIGURE 2. Lorenz curves for wealth.

The decreases in wealth concentration are quite large in the economy—an increase in θ to 0.1 causes a 43% decline in the Gini coefficient. However, because we recalibrate the economy, it is difficult to distinguish the direct effect of increasing θ from the required decline in β needed to match the average amount of wealth. As β declines, households have shorter effective horizons; as a result, their savings functions decrease in slope and therefore generate smaller stationary points. With a lower upper bound and a fixed lower bound on wealth, the Gini coefficient will naturally decline as more agents have sufficiently good luck to reach their endogenous upper bound. To remove this effect, we hold β fixed at the calibrated value for the benchmark and examine how the wealth distribution changes when θ changes. The results are presented in Table 2A; with fixed β , the level of wealth increases dramatically but the concentration of wealth still shrinks. However, the ergodic set does not shrink; rather, it increases in size and dispersion in wealth actually increases. It is curious to note that the net return on capital $r - \delta$ becomes negative when θ rises sufficiently high; for example, when $\theta = 5.0$ the net return to capital with fixed (β, μ, δ) is -0.0012 . Because wealth is still producing status, it is demanded despite its negative value as a savings vehicle.¹⁹

TABLE 2A. Moments of wealth distribution, fixed (β, μ, δ) , $\gamma = 0$ (wealth not a luxury good)

Models	Gini	Mean (k)	Std. (k)	Skew (k)	Kurt. (k)	Max (k)	Min (k)
$\theta = 0.1$	0.149	2.659	0.704	0.413	3.252	7.336	0.362
$\theta = 0.5$	0.080	5.655	0.800	0.279	3.127	10.222	2.526
$\theta = 1.0$	0.058	9.109	0.941	0.259	3.177	14.550	5.172
$\theta = 5.0$	0.033	25.642	1.521	0.167	3.106	34.149	18.759

TABLE 2B. Moments of consumption distribution, fixed (β, μ, δ) , $\gamma = 0$ (wealth not a luxury good)

Models	Gini	Mean (c)	Std. (c)	Skew (c)	Kurt. (c)
$\theta = 0.1$	0.034	0.519	0.032	-0.074	3.094
$\theta = 0.5$	0.030	0.592	0.032	0.048	3.061
$\theta = 1.0$	0.028	0.610	0.031	0.052	2.926
$\theta = 5.0$	0.021	0.435	0.016	0.156	2.934

TABLE 2C. Moments of hours distribution, fixed (β, μ, δ) , $\gamma = 0$ (wealth not a luxury good)

Models	Gini	Mean (h)	Std. (h)	Skew (h)	Kurt. (h)	Corr (h, y)
$\theta = 0.1$	0.143	0.344	0.087	-0.419	2.961	0.905
$\theta = 0.5$	0.112	0.399	0.079	-0.389	2.870	0.925
$\theta = 1.0$	0.090	0.455	0.073	-0.382	2.830	0.934
$\theta = 5.0$	0.036	0.690	0.044	-0.364	2.717	0.955

The intuition underlying the adjustments in (β, μ, δ) necessary to match the calibration targets across different values of θ is simple and therefore summarized briefly (a more complete discussion can be found in the working paper version). With higher θ , asset accumulation is larger by households for all values of $r - \delta$, requiring a lower value of β to hit the capital-output ratio and a higher value of μ to keep aggregate labor supply from rising.

3.2.2. Consumption. The flipside to smoothing wealth is that consumption may now be less insulated against fluctuations. To examine this possibility, we compute the cross-sectional distribution of consumption in the steady state. As θ increases, we observe the following facts from Table 1B:

1. Mean consumption falls—the substitution effect caused by increases in the demand for status is larger than the wealth effect generated by additional capital and the direct increase caused by falling β ;

2. The standard deviation of consumption rises—consumption becomes more exposed to income risk as the demand to smooth status induces smoother asset positions;
3. Skewness in consumption rises;
4. Kurtosis in consumption rises and then falls;
5. The Gini coefficient for consumption rises.

Thus, we see some evidence that consumption is being exposed to more risk. As before, we cannot easily make statements about the effect of θ on consumption, since it is contaminated by the required changes in β . Table 2B presents the distribution of consumption statistics when (β, μ) are held fixed; it is clear that the increase in the Gini coefficient is the result of recalibration. However, there is a nonmonotonic effect on the standard deviation of consumption: it initially rises and then falls. The eventual decline in the standard deviation is the result of the massive increases in wealth evident in Table 2A; with more wealth, agents are better able to self-insure against movements in their income. At low levels for θ , however, increases in θ have the effect of raising the standard deviation of consumption. Furthermore, the nonmonotonic behavior of the correlation between hours and productivity disappears. As θ rises, there is an increase in this correlation; the most productive agents begin to work more to accumulate additional wealth and status, and this shift produces an increase in aggregate productivity.

In this specification, there is a strong disincentive to be very poor; the marginal utility of status goes to infinity as wealth goes to zero. Unlike the standard model, all of our households will hold positive stocks of assets even when there is no possibility of drawing zero income in a given period (as in our model). We can easily see this from Table 1A, where the minimum wealth in the distribution rises significantly as θ increases from zero.²⁰ Thus, counterfactually the model with $\gamma = 0$ predicts zero consumers who have zero wealth, and preferences are not even defined for negative levels of wealth. Unfortunately for this preference specification, Budría et al. (2002) report 9.9% of all households have zero or negative wealth. Clearly, this model is incapable of reproducing this observation; we will therefore examine how nonhomotheticity in the preference for status affects the wealth distribution.

3.3. Luxury Goods I: $\gamma > 0$

We now consider Stone-Geary preferences over status; these preferences can allow households to consider negative wealth positions. Although we now define utility over negative wealth positions, we do not allow households to borrow; we make this assumption to maintain comparability across model specifications. As noted earlier, status will be a luxury good whenever $\gamma > 0$. This assumption would seem plausible given that membership in country clubs, philanthropic contributions, and other status-enhancing activities are strongly correlated with wealth.²¹ Carroll (2002) suggests that nonhomothetic preferences over wealth (in his formulation, bequests) can account for the portfolios chosen by the very wealthy, whereas

TABLE 3A. Moments of wealth distribution $\gamma = 2.0$ (wealth a luxury good)

Models	Gini	Mean (k)	Std. (k)	Skew (k)	Kurt. (k)	Max (k)	Min (k)
$\theta = 0.1$	0.271	1.919	0.926	0.458	3.177	8.057	0.000
$\theta = 0.5$	0.177	1.912	0.600	0.276	3.145	5.532	0.000
$\theta = 1.0$	0.149	1.909	0.504	0.257	3.136	4.811	0.000

TABLE 3B. Moments of consumption distribution $\gamma = 2.0$ (wealth a luxury good)

Models	Gini	Mean (c)	Std. (c)	Skew (c)	Kurt. (c)
$\theta = 0.1$	0.036	0.480	0.031	-0.146	3.774
$\theta = 0.5$	0.039	0.478	0.033	0.102	3.057
$\theta = 1.0$	0.042	0.477	0.035	0.147	3.048

TABLE 3C. Moments of hours distribution $\gamma = 2.0$ (wealth a luxury good)

Models	Gini	Mean (h)	Std. (h)	Skew (h)	Kurt. (h)	Corr (h, γ)
$\theta = 0.1$	0.165	0.327	0.096	-0.441	2.950	0.890
$\theta = 0.5$	0.152	0.327	0.088	-0.434	3.028	0.879
$\theta = 1.0$	0.146	0.327	0.085	-0.443	3.099	0.867

Reiter (2004) argues that it can help account for the savings behavior of the very wealthy.

When we compute our model with $\gamma = 2.0$, we see that the Gini coefficient on wealth is higher for every value of θ considered; Figure 3 shows the Lorenz curve for wealth for various different values of γ when $\theta = 0.1$, and shows clearly that the gap between the curve and the 45° line gets larger as γ rises. Additionally, Tables 3A–3C present the moments of wealth, consumption, and hours. It is clear that the ergodic set of capital stocks shrinks much more slowly than when $\gamma = 0$, and the lower bound remains at the borrowing limit instead of rising. Standard deviations are also smaller in this case. What γ is doing is muting the increased demand for capital by reducing the marginal utility of status for every agent, and this effect is larger the smaller the wealth of the agent. Further increases in γ (for example, to twice the average amount of capital) increase the Gini coefficient on wealth. Reiter (2004) uses a value that is equal to 30,000 times the capital/output ratio in his economy. When we solve this economy, we find that the Gini coefficient is exactly the same when $\theta = 0.1$ as when $\theta = 0.0$. Nonhomothetic preferences over status cannot increase the Gini coefficient, however, as further increases in γ have no effect on inequality.²² This result obtains in our economy because, with separable preferences, as $\gamma \rightarrow \infty$ the households behave identically to ones with $\theta = 0$; with an arbitrarily large constant in the utility function, status is unaffected




FIGURE 3. Lorenz curves for wealth, $\theta = 0.1$.

by wealth. $\gamma > 0$ has little effect on the distribution of hours as it primarily affects the wealthy and these households supply little labor.

3.4. Luxury Goods II: $\sigma < 1$

We next consider the effects of reducing the risk aversion over status gambles. Because it is intuitively implausible that status is a necessity good, we investigate only values with $\sigma < 1$; as noted earlier, this parametrization implies that status is a luxury good. Specifically, we explore whether the model can produce a higher Gini coefficient for wealth when $\sigma < 1$ than the version with $\theta = 0$; the effects of reducing σ on savings are very similar to those generated by increasing γ . With $\sigma = 0.5$ we obtain the results in Table 4; the concentration of wealth is higher the lower σ is. With less risk aversion over status, agents are more willing to permit the ergodic set of wealth to spread out, resulting in more concentration and higher measures of dispersion. But the direct effect of θ still produces large declines in wealth concentration. Figure 4 presents the Lorenz curves as σ varies from 0.01 to 0.5; even with near-risk-neutrality with respect to status, wealth Gini coefficients decline relative to the benchmark.²³

TABLE 4A. Moments of wealth distribution $\sigma = 0.5, \gamma = 0.0$ (wealth not a luxury good)

Models	Gini	Mean (k)	Std. (k)	Skew (k)	Kurt. (k)	Max (k)	Min (k)
$\theta = 0.1$	0.198	1.914	0.674	0.464	3.285	6.494	0.000
$\theta = 0.5$	0.134	1.906	0.455	0.389	3.236	4.691	0.362
$\theta = 1.0$	0.119	1.904	0.402	0.378	3.226	4.210	0.603

TABLE 4B. Moments of consumption distribution $\sigma = 0.5, \gamma = 0.0$ (wealth not a luxury good)

Models	Gini	Mean (c)	Std. (c)	Skew (c)	Kurt. (c)
$\theta = 0.1$	0.038	0.479	0.032	-0.061	3.153
$\theta = 0.5$	0.043	0.477	0.037	0.064	3.027
$\theta = 1.0$	0.045	0.476	0.038	0.098	3.014

TABLE 4C. Moments of hours distribution $\sigma = 0.5, \gamma = 0.0$ (wealth not a luxury good)

Models	Gini	Mean (h)	Std. (h)	Skew (h)	Kurt. (h)	Corr. (h, y)
$\theta = 0.1$	0.155	0.327	0.090	-0.447	3.029	0.882
$\theta = 0.5$	0.142	0.327	0.083	-0.468	3.189	0.857
$\theta = 1.0$	0.138	0.327	0.081	-0.482	3.272	0.845

3.5. Discussion

Why do we find that soc preferences reduce wealth concentration but De Nardi (2004) finds the opposite? The key difference in the models is that the preference for bequests in De Nardi (2004) is not uniform across agents. In her model, agents age stochastically; only retirees are subject to mortality risk and therefore only those agents directly value bequests. We can replicate this feature in our model easily by assuming that some agents have soc preferences and some do not. It turns out that this model can easily replicate the observed wealth Gini coefficient in the data because it looks very similar to the stochastic discount factor model of Krusell and Smith (1998). Agents with soc preferences demand a lot of assets and become very wealthy; with a concave production technology the presence of these agents drives down the equilibrium return, reducing the asset holdings of the households with standard preferences. Roughly speaking, these two groups act like the high and low discount factor types in Krusell and Smith (1998), respectively. For any fixed measure of soc households we can then generate the observed Gini coefficient merely by choosing an appropriate value of θ , which is akin to choosing the gap between discount factors. Thus, a version of our model with heterogeneous




FIGURE 4. Lorenz curves for wealth, $\theta = 0.1$.

preferences does not add anything new to the mix of potential resolutions of the wealth concentration puzzle, it simply relabels an existing one.²⁴

4. TAX POLICY

In this final substantive section of the paper, we consider some experiments designed to explore whether spirit of capitalism preferences have any substantive impact on the evaluation of certain types of taxes. Specifically, we consider the replacement of a progressive income tax with two different flat tax systems—an income tax system and a consumption tax system. Our flat income tax experiments are similar in spirit to those conducted by Castañeda, Díaz-Giménez, and Ríos-Rull (2003) and Conesa and Krueger (2006), who compute the welfare and distributional consequences of replacing a calibrated progressive tax system with a revenue-neutral flat tax.²⁵ Our interest is not in a careful measurement of the relative gain from reforming the tax system, but rather to assess whether spirit of capitalism preferences have consequences for these types of policy experiments.²⁶ We find that the predictions of the baseline model are robust to the introduction of observably equivalent soc specifications.

Following Gouveia and Strauss (1994), we approximate the existing income tax code using the function

$$\tau(i) = a_0 \left(i - (i^{-a_1} + a_2)^{-\frac{1}{a_1}} \right) + a_3 i, \quad (4.1)$$

with parameters (a_0, a_1, a_2, a_3) . The scale-invariant parameters of the progressive component are $a_0 = 0.258$ and $a_1 = 0.768$. The other parameters are chosen such that the progressive component finances 68 percent of government spending and the government budget constraint is balanced:

$$\int (\tau_c c + \tau(i)) \Gamma(k, y) = G, \quad (4.2)$$

where $\tau_c = 0.05$ is the fixed consumption tax. a_3 is meant to capture all the nonconsumption, nonincome taxes levied by the government in the data (tariffs, estate taxes, user fees, and a variety of other revenue sources); we implicitly assume that the incidence of this tax is proportional to income. In the calibrated equilibrium it turns out to be small quantitatively, so we do not consider this assumption to be of any importance.

The budget constraint of the household is now given by

$$(1 + \tau_c)c + k' = i + k - \tau(i), \quad (4.3)$$

where

$$i = (r - \delta)k + wyh, \quad (4.4)$$

is personal income. The goods market becomes

$$C + K' + G = Y + (1 - \delta)K. \quad (4.5)$$

The definition of equilibrium needs only to be modified to account for these two changes.

We consider only revenue neutral experiments here, so we hold fixed G after the tax reform.²⁷ We then solve endogenously for the transition path to the new steady state and evaluate welfare using the value function at the time the transition began. This function identifies the welfare change from the tax reform as a function of the initial state and can be converted into consumption units as

$$\phi(k, y) = \exp((1 - \beta)v^r(k, y) - v^0(k, y)) - 1,$$

where v^r is the transition value function and v^0 is the lifetime utility an agent would experience if no tax reform was undertaken. We can obtain an aggregate welfare measure by integrating $\phi(k, y)$ with respect to $\Gamma^0(k, y)$, the invariant distribution of the economy with progressive taxation, which we denote Φ :

$$\Phi = \int \phi(k, y) \Gamma^0(k, y);$$

TABLE 5. Income tax experiments

Experiment	(θ, γ)	ΔMean (k)	ΔGini (k)	ΔMean (c)	ΔGini (c)	ΔMean (h)	Φ
Constant G	(0.0, 0.0)	7.1%	10.6%	6.4%	18.6%	3.9%	0.0137
	(0.1, 50.0)	6.1%	11.3%	5.2%	19.6%	4.0%	0.0128
	(0.2, 100.0)						0.0159
Consumption Tax Experiments							
	(θ, γ)	ΔMean (k)	ΔGini (k)	ΔMean (c)	ΔGini (c)	ΔMean (h)	Φ
Constant G	(0.0, 0.0)	33.7%	10.0%	15.5%	12.1%	5.8%	0.0263
	(0.1, 50.0)	34.0%	10.3%	15.2%	12.2%	5.9%	0.0257
	(0.2, 100.0)						0.0289

note that this measure is not the usual one, which typically takes the difference of the expected utility and then converts it to consumption units, but is instead consistent with the individual-specific measures computed using the transition path.²⁸ In order to make the experiments as informative as possible, we would prefer to choose parameterizations that imply the same initial wealth distributions independent of (θ, γ) . For the reasons noted earlier, we are unable to do this entirely, so we instead choose a parametrization for the soc case that produces nearly the same Gini coefficient for wealth as the benchmark case but is not too large a departure: $\theta = 0.1$ and $\gamma = 50.0$. The Lorenz curves for wealth for the two economies lie essentially on top of each other, and the cumulative distributions of wealth are nearly impossible to visually distinguish. We also consider the case $\theta = 0.2$ and $\gamma = 100.0$; this parametrization again produces the same wealth distribution, but has a stronger demand for status for agents who are sufficiently wealthy.²⁹

Table 5 and Figures 5–8 present our results from considering the two types of tax reforms. For standard preferences, Table 5 shows that there is a significantly larger welfare gain from reforming to a consumption tax over an income tax, holding government spending fixed. In Figure 5 one can see that the consumption tax reform induces an FSD shift in the distribution of wealth, so that lifetime expected utility is higher for any increasing value function (although not shown, the FSD shift holds for all values of y).³⁰ The welfare gain is large, on average over 2.6% of individual consumption, compared to only 1.3% for the income tax reform. In contrast, the income tax reform has a smaller effect on the distribution of wealth; unfortunately, the shift is not ranked by any stochastic dominance ordering. The picture is very similar in Figure 6, which considers the soc case $(\theta, \gamma) = (0.1, 50.0)$; the effects on the wealth distribution are essentially unchanged and the welfare changes are only slightly smaller on average. However, the welfare gains are larger when $(\theta, \gamma) = (0.2, 100.0)$ for both types of tax reform.




FIGURE 5. CDF for wealth, benchmark.

In Figure 7, we present the welfare gains as functions of the initial wealth holdings of the individuals for the mean level of initial productivity. Qualitatively, the pictures show that consumption tax reforms benefit the initially poor and harm the initially wealthy, while income tax reforms benefit all wealth levels.³¹ Thus, an income tax change is a Pareto-improvement while a consumption tax generates larger average welfare gains.³² All gains are increasing in initial productivity. Some individuals can gain quite a significant amount from a consumption tax reform, notably those households who are initially asset-poor but have currently high productivity. Consumption tax reform raises wages more than income tax reform does. Because initially asset-constrained agents care mostly about wages, their gains are large, particularly for those agents who are currently highly productive and therefore expect to still be productive in the future; wealthier agents, who care less about wages and more about returns, lose under the consumption tax reform because returns fall quite a lot.

As we noted in the Introduction, the results regarding the relative efficiency of income versus. consumption tax reform seem to be new. They do, however, echo existing results in the literature regarding the welfare gains generated by taxing capital in models with incomplete asset markets. When $\theta = 0$, Aiyagari (1995) shows that an economy with inelastic labor supply and incomplete markets

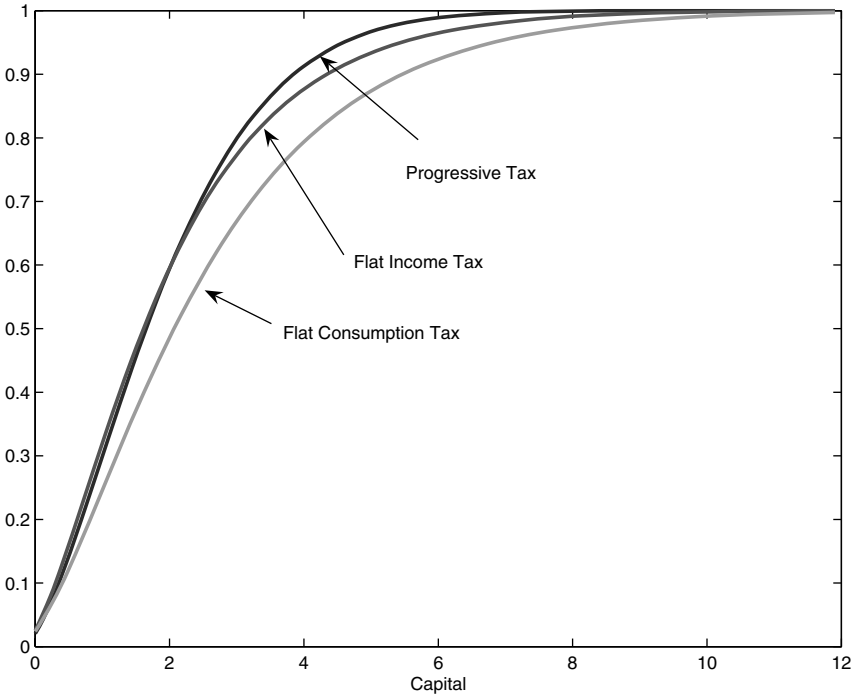


FIGURE 6. CDF for wealth ($\theta = 0.1$, $\gamma = 50.0$).

suffers from overaccumulation of capital; this overaccumulation can be rectified by taxing capital and reducing it to the level of the modified Golden rule where $r - \delta = \beta^{-1} - 1$.³³ Because the consumption tax reform does not tax capital, our economy may suffer from the same overaccumulation, limiting the appeal for agents whose initial wealth is high.³⁴ In equilibrium there are not enough of these agents to make the consumption tax unattractive on average. Quantitatively the same argument goes through when $\theta > 0$, even though the consumption tax reform indirectly taxes capital.

Figure 8 displays the transitional dynamics of the interest rate across the four cases—standard and soc preferences, consumption and income tax reforms. The dynamics are very similar across preference specifications—the interest rate declines monotonically, essentially reaching its new steady state value in 75 years; wages move in the opposite direction, rising for each type of tax reform. The decline in the interest rate is larger for soc preferences for both types of tax reform, so the rise in the wage is also larger. The extra decline is caused by the increased demand for assets generated by positive values of θ . Both tax reforms shift the aggregate asset supply curve to the right, leading to the decline in returns and the rise in the capital/labor ratio.




FIGURE 7. Mean initial productivity.

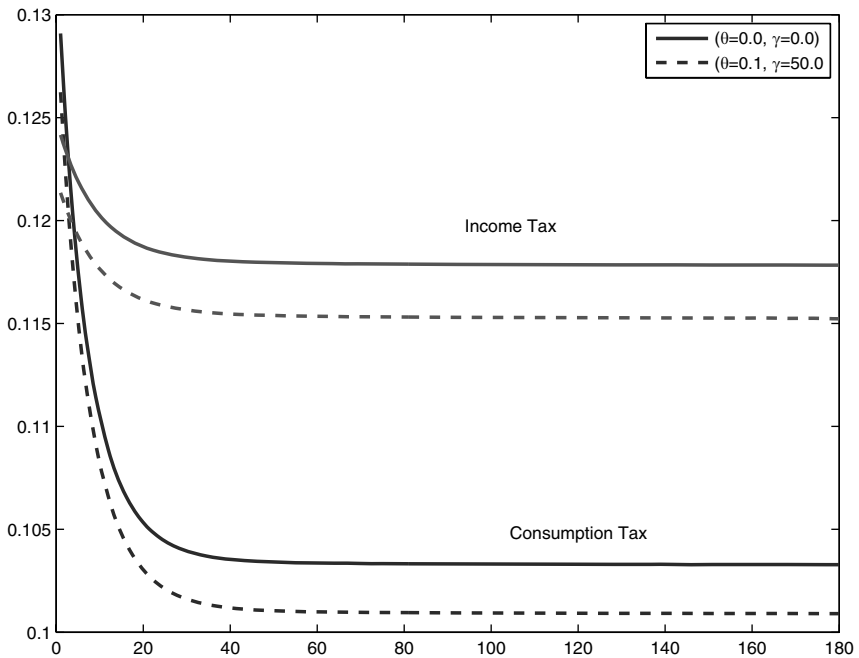


FIGURE 8. Transitional dynamics for r .

5. CONCLUSION

There may be alternative approaches to introducing status as a good that have better implications for the concentration puzzle. Specifically, it is a reasonable hypothesis that status is affected by two variables: wealth and some particular kind of consumption spending (such as charitable giving or conspicuous consumption purchases). If these components are complementary, we could find large differences across agents in terms of their spending on status, particularly if the household objective is not concave; a similar impact can be seen in the literature on culture consumption [such as Krusell and Stavlöt (2005)], in which agents with different wealth levels purchase very different baskets of goods. Equilibria for that model would be considerably more difficult to compute, since it likely would possess many stationary distributions; given the non-convexity of the decision problem faced by households, the invariant distribution would depend on the initial distribution, so generality would be hard to achieve. We are currently exploring the computational tools needed to solve this extension.³⁵

We think it advisable to consider the asset pricing implications of our model. In Krusell and Smith (1997), the asset pricing behavior of their benchmark model with exogenous labor supply and aggregate shocks was shown to be quite poor. The essence of the problem is that only a small fraction of agents price bonds in their economy, and these agents are quite well insured. As a result, their marginal rates of substitution do not vary much in equilibrium, creating very little improvement in the failures of the complete markets model; this anomaly is not resolved by the introduction of stochastic discount factors.³⁶ Introducing spirit of capitalism preferences could potentially alter the nature of asset pricing within their model, because it implies that the elasticity of intertemporal substitution depends on $\frac{c}{k+\gamma}$; depending on the particular relationship between idiosyncratic and aggregate risk, this value could move countercyclically, a feature that has been shown to have some importance for asset pricing.³⁷ This change might be particularly pronounced if status is a luxury good, as it could imply sharp behavioral differences for wealthy versus poor households; this break would be helpful for resolving the equity premium puzzle as the poor will generally determine the risk-free rate, whereas the rich will determine the equity return.

NOTES

1. This preference structure also could be motivated by appeals to home production technologies, in which certain components of wealth are used to produce home consumption goods. However, a critical difference in that setup and one with spirit of capitalism is that current wealth confers utility without any opportunity cost.

2. The improvement in asset pricing may be illusory because of the failure of the papers to impose all the restrictions implied by the model; see Lettau (1997) for a discussion.

3. This approach is advocated by Castañeda, Díaz-Giménez, and Ríos-Rull (2003).

4. The fact that the efficient allocation may have more capital is surprising given that precautionary savings yields a return on capital below the time rate of preference in this model. A similar effect is noted in Pijoan-Mas (2006) using a model with elastic labor supply.

5. We do not solve the social planning problem in this paper, although it would be an interesting avenue for future research.

6. Francis (2005a) finds that soc preferences can increase wealth concentration in a life-cycle model. Francis (2005b) uses the same model with an entrepreneurial choice to obtain similar results, including higher wealth concentration among entrepreneurs than workers. There are two assumptions in her model that seem important—exogenous prices and inoperative bequest motives. We plan to explore explicitly the role of these assumptions in future research.

7. Krusell, Quadrini, and Ríos-Rull (1996) contains a discussion of the large public finance literature that finds consumption taxes dominate income taxes. This literature does not investigate the optimality of consumption taxation in the model used here, however.

8. This model has become the standard vehicle for exploring inequality in the quantitative macroeconomic literature.

9. The working paper version contains a discussion of the Engel curves for the soc preference setup. For brevity, this discussion has been omitted.

10. This assumption does not play an important role in determining the effects of soc on the concentration of wealth. Assuming that agents face liquidity constraints but a full set of Arrow securities are marketed implies the same effects as the incomplete market economy does, but the overall amount of concentration falls dramatically.

11. In a previous version of the paper we explored the business cycle dynamics of this model using tools developed in Young (2005). The impact of wealth in the utility function turned out to be trivial. In the Conclusion, we discuss extending our model to consider asset pricing, an extension for which aggregate shocks are obviously critical.

12. The convergence result depends on the concavity of the consumption function. We can verify numerically that this limit exists, but we have been unable to prove it. The resulting bound on $r - \delta$ was verified across a wide range of parameter settings.

13. It is not possible to use β to match the wealth/output ratio and θ to match the risk-free return because the wealth/output ratio is the only argument of the marginal product of capital. In previous versions of the model, we found that risk aversion over consumption was not important for our results, provided it does not imply status is a necessity good.

14. The wage process is estimated using the Panel Study of Income Dynamics (PSID), whereas the facts about the wealth distribution are taken from the Survey of Consumer Finances (SCF). The second dataset has no time series dimension but it oversamples the wealthy, making it unsuitable for estimating shock processes but better at describing cross-sectional facts. The absence of the very wealthy from the PSID means that no agents draw very high wages and as a result never become extremely wealthy.

15. The small irregular spikes occur at points where households in productivity state y hit the zero constraint on hours. We did not include the distributions for the higher values of θ because they are simply more pronounced versions of the $\theta = 0.5$ case and scaling starts to become an issue.

16. It is important to note here that our invariant distribution is not constructed using a finite simulation, meaning that sample size concerns are not an issue. See the technical appendix to Young (2004) for further explanation.

17. We confirmed this fact by computing a version of the model with $\theta = 10,000$, which resulted in all the mass of the invariant distribution being located over two points in the grid, those that bracket the desired aggregate stock of capital. Adding additional grid points would produce a better estimate of the limit point, but there seems to be no reason to do so.

18. When $\theta = 10,000.0$, the case that leads to the complete collapse of the wealth distribution, the correlation between h and y is 0.8758.

19. If we allowed households to store consumption goods, the net return on capital could not go below zero in equilibrium as no capital would be held, although it is reasonable to think that stored consumption goods would not confer status.

20. Again, we wish to point out that this increase in the observed lower bound for wealth is not the result of simulation error, as our method for constructing the invariant distribution does not use finite simulations.

21. This correlation would follow directly from Veblen’s notion of conspicuous consumption; large philanthropic gifts confer status both because they are large (which requires high wealth) and because they are very public.

22. The cutoff value of γ (that is, where further increases have no effect on the wealth distribution) depends on the endogenous choice of the observed upper bound for k (which of course is determined by γ). As θ rises, this cutoff value for γ falls.

23. Letting σ be larger than 1 produces a concave expansion path for consumption and leads to even less wealth concentration; this effect is a result of risk aversion over wealth that tends to reduce the support of the invariant distribution.

24. That is not necessarily to say that the policy implications of the two models would be the same.

25. The consequences of taxation for inequality is also investigated by Ventura (1999), Eicher, Riera Prunera, and Turnovsky (2003), and Díaz-Giménez and Pijoan-Mas (2006).

26. Our paper is not the only one in which nonstandard preferences have been used to explore the effects of progressive income taxation; other examples include Boskin and Sheshiski (1978) and Corneo (2002).

27. We do not attempt to search for the optimal tax code, as is done in Conesa and Krueger (2006).

28. The alternative method is to compute

$$V^0 = \int v^0(k, y)\Gamma_0(k, y)$$

as no-reform utility and

$$V^r = \int v^r(k, y)\Gamma_0(k, y)$$

as reform utility, and then convert to consumption units:

$$\phi = \exp((1 - \beta)(V^r - V^0)) - 1.$$

Aiyagari and McGrattan (1998) contains a discussion of interpretations of these welfare measures.

29. The working paper version contains a discussion of a two-period model that highlights the tradeoffs generated by taxes with soc preferences. For brevity, this discussion has been omitted.

30. Definitions of stochastic dominance criteria are provided in Rothschild and Stiglitz (1970) and Hadar and Russell (1969). Because these pictures only present the stationary distribution, the comments about FSD shifts are not relevant for the welfare calculations relating to the transition path.

31. Qualitatively similar pictures can be shown for all levels of productivity; we omit these in the interest of brevity.

32. The intermediate case $(\theta, \gamma) = (0.1, 50)$ is impossible to visually distinguish from $(\theta, \gamma) = (0, 0)$ and is omitted; however, it is clear from the average welfare calculation that the curves have shifted downward.

33. A more detailed model that also displays the efficiency of capital taxation can be found in Conesa, Kitao, and Krueger (2007), which includes overlapping generations of finitely-lived agents and separate progressive tax functions for capital and labor income.

34. We say may because of the previously-mentioned results in Pijoan-Mas (2006); because the distribution of hours changes in the presence of idiosyncratic wage risk, the economy may end up with less capital rather than more. The bound on the interest rate is unchanged. Domeij and Heathcote (2004) present some results on how the preference for capital and labor income taxes are distributed in a model similar to ours.

35. These tools would have value beyond our application, as they would apply to any economic problem where the objective function is not globally concave.

36. Krueger and Lustig (2007) discuss the reasons underlying the failure of models with idiosyncratic risk to resolve the equity premium puzzle; basically the representative agent is still pricing assets.

37. Boileau and Brown (2007) study a representative-agent business cycle model with soc and complete markets, finding many counterfactual implications.

REFERENCES

- Aiyagari, S. Rao (1994) Uninsured idiosyncratic risk and aggregate saving. *Quarterly Journal of Economics* 109, 659–684.
- Aiyagari, S. Rao (1995) Optimal capital income taxation with incomplete markets, borrowing constraints, and constant discounting. *Journal of Political Economy* 103, 1158–1175.
- Aiyagari, S. Rao and Fernando Alvarez (2001) Efficient Dynamic Monitoring of Unemployment Insurance Claims. Mimeo, University of Chicago.
- Aiyagari, S. Rao and Ellen R. McGrattan (1998) The optimum quantity of debt. *Journal of Monetary Economics* 42, 447–469.
- Bakshi, Gurdip S. and Zhiwu Chen (1996) The spirit of capitalism and stock-market prices. *American Economic Review* 86, 133–157.
- Boileau, Martin and Rebecca W. Brown (2007) The spirit of capitalism, asset returns, and the business cycle. *Macroeconomic Dynamics* 11, 214–230.
- Boskin, Michael J. and Eytan Sheshinski (1978) Optimal redistributive taxation when individual welfare depends on relative income. *Quarterly Journal of Economics* 92, 589–601.
- Budría Rodríguez, Santiago, Javier Díaz-Giménez, Vincenzo Quadrini, and José-Víctor Ríos-Rull (2002) Updated facts on the U.S. distributions of earnings, income, and wealth. *Federal Reserve Bank of Minneapolis Quarterly Review* 26, 2–35.
- Carroll, Christopher D. (2002) Portfolios of the rich. In Luigi Guiso, Michael Haliassos, and Tullio Jappelli (eds.), *Household Portfolios: Theory and Evidence*. Cambridge, MA: MIT Press.
- Carroll, Christopher D. (2004) Theoretical Foundations of Buffer Stock Savings. NBER Working Paper 10867.
- Castañeda, Ana, Javier Díaz-Giménez, and José-Víctor Ríos-Rull (2003) Accounting for wealth and income inequality. *Journal of Political Economy* 111, 818–857.
- Chue, Timothy K. (2004) The Spirit of Capitalism and International Risk Sharing. Econometric Society 2004 Far Eastern Meetings Working Paper 589.
- Conesa, Juan Carlos, Sagiri Kitao, and Dirk Krueger (2007) Taxing Capital? Not a Bad Idea After All! Mimeo, Goethe University.
- Conesa, Juan Carlos and Dirk Krueger (2006) On the optimal progressivity of the income tax code. *Journal of Monetary Economics* 53, 1425–1450.
- Corneo, Giacomo (2002) The efficient side of progressive income taxation. *European Economic Review* 46, 1359–1368.
- Dávila, Julio, Jay H. Hong, Per Krusell, and José-Víctor Ríos-Rull (2006) Constrained-Efficient Allocations in the Growth Model with Idiosyncratic Shocks. Mimeo, University of Pennsylvania.
- Díaz-Giménez, Javier and Josep Pijoan-Mas (2006) Flat Tax Reforms in the U.S.: A Boon for the Income Poor. CEPR Discussion Paper 5812.
- De Nardi, Mariacristina (2004) Wealth inequality and intergenerational links. *Review of Economic Studies* 71, 743–768.
- Domeij, David and Jonathan Heathcote (2004) On the distributional effects of reducing capital taxes. *International Economic Review* 45, 523–554.
- Duesenberry, James S. (1949) *Income, Saving, and the Theory of Consumer Behavior*. Cambridge, MA: Harvard University Press.

- Eicher, Theo, Maria Carmen Riera Prunera, and Stephen J. Turnovsky (2003) Tax policy and inequality: Empirical and theoretical implications. In Theo Eicher and Stephen J. Turnovsky (eds.), *Growth and Inequality: Issues and Policy Implications*. Cambridge, MA: MIT Press.
- Floden, Martin and Jesper Lindé (2001) Idiosyncratic risk in the United States and Sweden: Is there a role for government insurance? *Review of Economic Dynamics* 4, 406–437.
- Francis, Johanna L. (2005a) Life Cycle Saving and the Capitalist Spirit. Mimeo, Johns Hopkins University, Baltimore, MD.
- Francis, Johanna L. (2005b) Entrepreneurs and the Capitalist Spirit. Mimeo, Johns Hopkins University, Baltimore, MD.
- Gong, Liutang and Heng-fu Zou (2002) Direct preferences for wealth, the risk premium puzzle, growth, and policy effectiveness. *Journal of Economic Dynamics and Control* 26, 247–270.
- Gouveia, Miguel and Robert P. Strauss (1994) Effective federal individual tax functions: An exploratory empirical analysis. *National Tax Journal* 47, 317–339.
- Hadar, Josef and William R. Russell (1969) Rules for ordering uncertain prospects. *American Economic Review* 49, 25–34.
- Kenc, Turalay and Sel Diboğlu (2007) The spirit of capitalism, asset pricing and growth in a small open economy. *Journal of International Money and Finance* 26(8), 1378–1402.
- Kongsamut, Piyabha, Sergio Rebelo, and Danyang Xie (2001) Beyond balanced growth. *Review of Economic Studies* 68, 869–882.
- Krueger, Dirk and Hanno Lustig (2007) When Is Market Incompleteness Irrelevant for the Price of Aggregate Risk? Mimeo, University of Pennsylvania and UCLA.
- Krusell, Per, Vincenzo Quadrini, and José-Víctor Ríos-Rull (1996) Are consumption taxes really better than income taxes? *Journal of Monetary Economics* 37, 475–503.
- 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.
- Krusell, Per and Anthony A. Smith Jr. (2002) Revisiting the Welfare Effects of Eliminating Business Cycles. Mimeo, Princeton University and Yale University.
- Krusell, Per and Ulrika Stavlöt (2005) On the Origins of Diversity of Culture Consumption. Mimeo, Princeton University.
- Lettau, Martin (1997) Comment on “The Spirit of Capitalism and Stock Market Prices” by G. S. Bakshi and Z. Chen (AER, 1996). Tilburg University, Center for Economic Research Discussion Paper 49.
- Pjioan-Mas, Josep (2006) Precautionary savings or working longer hours? *Review of Economic Dynamics* 9, 326–352.
- Reiter, Michael (2004) Do the Rich Save Too Much? How to Explain the Top Tail of the Wealth Distribution. Mimeo, Universitat Pompeu Fabra.
- Robson, Arthur J. (1992) Status, the distribution of wealth, private and social attitude to risk. *Econometrica* 60, 837–857.
- Rothschild, Michael and Joseph E. Stiglitz (1970) Increasing risk I: A Definition. *Journal of Economic Theory* 2, 225–243.
- Smith, William T. (2001) How does the spirit of capitalism affect stock market prices? *Review of Financial Studies* 14, 1215–1232.
- Tauchen, George (1986) Finite state Markov-chain approximations to univariate and vector autoregressions. *Economics Letters* 20(2), 177–181.
- Ventura, Gustavo (1999) Flat tax reform: A quantitative exploration. *Journal of Economic Dynamics and Control* 23, 1425–1458.
- Young, Eric R. (2004) Unemployment insurance and capital accumulation. *Journal of Monetary Economics* 51, 1683–1710.
- Young, Eric R. (2005) Approximate Aggregation. Mimeo, University of Virginia.
- Zou, Heng-fu (1994) The spirit of capitalism and long-run growth. *European Journal of Political Economy* 10, 279–293.

APPENDIX

In this appendix, we show that for the class of utility functions $u(c, s, l) = \frac{[c(s+\gamma)^\theta l^\mu]^{1-\sigma}}{1-\sigma}$, status is a luxury good if $\theta < 1$ and $\gamma > 0$. Define

$$\eta_c = \frac{\frac{du_c(c, s, l)}{dc}}{\frac{u_c(c, s, l)}{c}} = -\theta - (1 - \theta)\sigma$$

$$\eta_s = \frac{\frac{du_s(c, s, l)}{ds}}{\frac{u_s(c, s, l)}{s}} = (-\theta\sigma + \theta - 1) \frac{s}{s + \gamma}.$$
(A.1)

With $\gamma > 0$ and $\theta < 1.0$ (with $\sigma \geq 1$) we have $\eta_s > \eta_c$; that is, the marginal utility of status declines less with wealth than the marginal utility of consumption does. In other words, as wealth rises so does the fraction of current utility derived from status.

We next compute the EIS for this utility function. Let status be simply given by current capital. The Euler equation can be written

$$u_c(c_t, k_t, l_t) = \beta E_t [u_c(c_{t+1}, k_{t+1}, l_{t+1})(1 + r_{t+1}) + u_s(c_{t+1}, k_{t+1}, l_{t+1})]. \quad (\text{A.2})$$

Using the functional form

$$u(c, s, l) = \frac{(c(k + \gamma)^\theta l^\mu)^{1-\sigma} - 1}{1 - \sigma},$$

the Euler equation becomes

$$c_t^{-\sigma} (k_t + \gamma)^{\theta(1-\sigma)} l_t^{\mu(1-\sigma)} = \beta E_t \left[\begin{array}{c} c_{t+1}^{-\sigma} (k_{t+1} + \gamma)^{\theta(1-\sigma)} l_{t+1}^{\mu(1-\sigma)} (1 + r_{t+1}) + \\ \theta c_{t+1}^{1-\sigma} (k_{t+1} + \gamma)^{\theta(1-\sigma)-1} l_{t+1}^{\mu(1-\sigma)} \end{array} \right]. \quad (\text{A.3})$$

Along some steady state path we have no uncertainty, so this becomes

$$1 = \beta \left[\begin{array}{c} \left(\frac{c_{t+1}}{c_t} \right)^{-\sigma} \left(\frac{k_{t+1} + \gamma}{k_t + \gamma} \right)^{\theta(1-\sigma)} (1 + r_{t+1}) + \\ \theta \left(\frac{c_{t+1}}{c_t} \right)^{-\sigma} \left(\frac{k_{t+1} + \gamma}{k_t + \gamma} \right)^{\theta(1-\sigma)-1} \frac{c_{t+1}}{k_{t+1} + \gamma} \end{array} \right], \quad (\text{A.4})$$

because l_t is constant in the steady-state path. Following Kongsamut, Rebelo, and Xie (2001), we define a Generalized Balanced-Growth path as one that implies a constant real interest rate r ; we leave aside the issue of whether the model possesses such a growth path. Constant r requires that

$$g = \frac{c_{t+1}}{c_t} = \frac{k_{t+1} + \gamma}{k_t + \gamma}. \quad (\text{A.5})$$

Note that, if $\gamma > 0$, we have that

$$g_s = \frac{k_{t+1}}{k_t} > \frac{c_{t+1}}{c_t},$$

because

$$\frac{d}{d\gamma} \left[\frac{k_{t+1} + \gamma}{k_t + \gamma} \right]_{\gamma=0} = \frac{k_t - k_{t+1}}{(k_t + \gamma)^2} < 0. \quad (\text{A.6})$$

Using this result, we have

$$1 = \beta \left[g^{-\sigma} (1+r) + \theta g^{-\sigma} \frac{c_{t+1}}{s_{t+1} + \gamma} \right]. \quad (\text{A.7})$$

The appearance of the additional term is what differentiates this model from the standard one. Rearranging, we obtain

$$g^\sigma = \beta \left[1 + r + \theta \frac{c}{k + \gamma} \right]. \quad (\text{A.8})$$

Taking logs, we obtain

$$\sigma \log(g) = \log(\beta) + \log(1+r) + \log \left(1 + \theta \frac{c}{k + \gamma} \frac{1}{1+r} \right).$$

Thus, there is a wedge between the EIS and $1/\sigma$; the size of the wedge is directly related to θ . For small enough values of θ and r , we obtain

$$\sigma(g - 1) \approx \log(\beta) + r + \theta \frac{c}{k + \gamma} \frac{1}{1+r},$$

or

$$\frac{dg}{dr} \approx \frac{1}{\sigma} - \frac{\theta}{\sigma} \frac{c}{k + \gamma} \left(\frac{1}{1+r} \right)^2 < \frac{1}{\sigma}. \quad (\text{A.9})$$

That is, the EIS is smaller than the standard model with $\theta = 0$. Increases in γ increase the EIS by decreasing the size of the wedge term, and as $\gamma \rightarrow \infty$ the EIS approaches the standard value $1/\sigma$.