Asset Pricing and the One Percent*

Alexis Akira Toda† Kieran Walsh‡

This version: December 30, 2014

Abstract

We find that when the income share of the top 1% income earners in the U.S. rises above trend by one percentage point, subsequent one year market excess returns decline on average by 5.6%. This negative relation remains strong and significant even when controlling for classic return predictors such as the price-dividend and the consumption-wealth ratios. To explain this stylized fact, we build a general equilibrium asset pricing model with heterogeneity in wealth and risk aversion across agents. Our model admits a testable moment condition and a novel two factor covariance pricing formula, where one factor is inequality. Intuitively, when wealth shifts into the hands of rich and risk tolerant agents, average risk aversion falls, pushing down the risk premium. Our model is broadly consistent with data and provides a novel positive explanation of both market excess returns over time and the cross section of returns across stocks.

Keywords: equity premium; heterogeneous risk aversion; return prediction; wealth distribution.

JEL codes: D31, D53, D58, G12, G17.

1 Introduction

Does the wealth distribution matter for asset pricing? Common sense tells us that it does: as the rich get richer, they buy risky assets and drive up prices. Indeed, over a century ago prior to the advent of modern mathematical finance, Fisher (1910) argued that there is an intimate relationship between prices, the heterogeneity of agents in the economy, and booms and busts. He contrasted (p. 175) the "enterpriser-borrower" with the "creditor, the salaried man, or the laborer," emphasizing that the former class of society accelerates fluctuations in prices and production. Central to his theories of fluctuations were differences in preferences and wealth across people.

Following the seminal work of Lucas (1978), however, the "representative agent" consumption-based asset pricing models—which seem to allow no role

*We benefited from comments by Dan Cao, Vasco Carvalho, Peter Debaere, Graham Elliott, John Geanakoplos, Jim Hamilton, Gordon Hanson, Fumio Hayashi, Toshiki Honda, Larry Schmidt, and seminar participants at Cambridge-INET, Carleton, Federal Reserve Board of Governors, Hitotsubashi ICS, Kyoto, Tokyo, UCSD, UVA, Vassar College, Yale, Yokohama National University, 2014 Australasian Finance and Banking Conference, and 2014 Northern Finance Association Conference.

†Department of Economics, University of California San Diego. Email: atoda@ucsd.edu.

‡Darden School of Business, University of Virginia. Email: walshk@darden.virginia.edu.

for agent heterogeneity—have dominated the literature, at least until recently. Yet agent heterogeneity may (and is likely to) matter even if a representative agent exists: unless agents have very specific preferences that admit the Gorman (1953) aggregation (a knife-edge case, which is unlikely to hold in reality), the preferences of the representative agent (in the sense of Constantinides (1982)) will in general depend on the wealth distribution, as pointed out by Gollier (2001). Indeed, even with complete markets, the preferences of the representative agent are typically nonstandard when individual utilities do not reside within quite particular classes.

To see the intuition as to why the wealth distribution affects asset pricing, consider an economy consisting of people with different attitudes towards risk or beliefs about future dividends. In this economy, equilibrium risk premiums and prices balance the agents' preferences and beliefs. If wealth shifts into the hands of the optimistic or less risk averse, for markets to clear, prices of risky assets must rise and risk premiums must fall to counterbalance the new demand of these agents. In this paper, we establish both the theoretical and empirical links between income/wealth inequality and asset prices.

This paper has two main contributions. First, we build a simple general equilibrium model of asset prices with heterogeneous agents and derive testable implications linking asset returns and inequality across risk aversion types. We show that the stochastic discount factor depends both on market returns and average risk tolerance, which depends on the wealth distribution. Although the connection between the heterogeneity of agents' risk aversion and asset prices has been recognized at least since Dumas (1989) and recently emphasized by Gârleanu and Panageas (2012), the literature presents few testable implications that can be easily examined in financial data. We thus provide a link allowing us to subject the theory to empirical scrutiny. The model also implies a new covariance pricing formula regarding the cross section of returns across assets: the average return of a stock depends *positively* on its correlation with the market and *negatively* on its correlation with the wealth share of risk tolerant agents. In short, the wealth distribution determines an asset pricing factor: average risk aversion across agents. When average risk aversion is low, average marginal utility is high. Therefore, assets correlated with top wealth shares (and thus average risk tolerance) command relatively low risk premiums.

We also illustrate the effect of the wealth distribution on asset prices with numerical examples. Two agents inhabit the model and trade a riskless bond in zero net supply and a risky asset in positive supply. We consider two examples, one in which agents have constant but heterogeneous relative risk aversion preferences, and another in which agents have identical but decreasing relative risk aversion preferences. Note that with declining relative risk aversion the wealthy are endogenously less risk averse. We perform comparative statics with respect to the initial endowment share of the more risk tolerant agent and find an inverse relationship between his income and the subsequent equity premium. In line with intuition, as the risk tolerant rich get richer, they buy risky assets, increasing their relative price. Subsequent excess returns thus fall.

Second, we empirically explore the theoretical predictions. We find that when the income share of the top 1% income earners in the U.S. is above trend, the subsequent one and five year U.S. stock market equity premiums are below average. That is, current inequality appears to forecast the subsequent risk premium of the U.S. stock market. Many heterogeneous agent general equilibrium

models in both macroeconomics and finance predict a relationship between the concentration of income and asset prices (see Section 1.1). We thus provide empirical support for a literature which has been subject to relatively little direct testing. Furthermore, the patterns we uncover are intuitive. In short, if one believes top earners are all else equal more willing to trade risk for return, then it should not be surprising that in the data asset returns suffer as the rich get richer.

More specifically, we employ regression analysis to establish the correlation between inequality and returns. Regressions of the year t to year $t+1$ excess return on the year t top 1% income share indicate a strong and significant negative correlation: when the top 1% income share rises above trend by one percentage point, subsequent one year market excess returns decline on average by 5.6%. This relation is strongly statistically significant and admits an R-squared of 9%. We show that estimated top wealth share series and the .1% income share also negatively predict subsequent returns. Furthermore, the top 1% income share predicts asset returns even after we control for some classic return forecasters such as the price-dividend ratio (Shiller, 1981) and the consumption-wealth ratio (Lettau and Ludvigson, 2001). It appears that the top 1% income share is not simply a proxy for the relative price level, which previous research shows correlates with subsequent returns. This is perhaps surprising because one imagines the rich being disproportionately exposed to stock price fluctuations. Our findings are also robust to the exclusion of capital gains in the income share series, to the inclusion of macro control variables, and to a large variety of detrending methods. Top marginal income tax rates are correlated with the 1% share but seemingly uncorrelated with subsequent excess returns. Using the top tax rate as an instrument magnifies the estimated impact of inequality on returns (at both the one and five year horizons) but also increases the standard error of the estimate.

We also empirically investigate the key moment conditions of the model. We structurally estimate the model's moment conditions to explore the extent to which the 6 Fama-French portfolios (sorted by size and book-to-market ratio) are explained by their relationship with market returns and inequality between two risk aversion types. As conjectured, we find that the rich are more risk tolerant than are the poor. We fail to reject our models and show that they outperform their homogeneous agent (standard CAPM) counterpart with respect to the Fama-French portfolios. We provide estimates of average risk tolerance over time in the U.S. and argue that its fluctuations are qualitatively in line with Irving Fisher's narratives.

Because any asset pricing model is inherently misspecified, we supplement the GMM exercise by the misspecification-robust two-pass estimation of Kan et al. (2013). In the first-pass time series regression, the income share coefficient is significant for three of six Fama-French portfolios (the "small" ones). Moreover, in line with our model, the income share coefficients are inversely related to the average portfolio returns: high return portfolios like "small size, high book-market" are negatively correlated with the top 1% share, and low return portfolios like "big size, low book-market" are positively correlated with the top 1% share. To reiterate, our explanation for this pattern is very simple and relies only on the 1% being more risk tolerant: as the rich get richer, the economy's risk tolerance increases and marginal utility rises. Thus, assets negatively correlated with the 1% have low payouts in high marginal utility states and thus

command a high risk premium. In the second-pass regression, our two factor model performs similarly to the Fama-French three factor model. And, according to the misspecification-robust estimation, the risk premium on the top 1% income share is negative and significant.

The rest of the paper is organized as follows. After discussing the literature, Section 2 introduces two simple models in which heterogeneity matters for asset pricing. We derive testable implications and provide numerical examples. Section 3 establishes the empirical link between inequality and subsequent asset returns. Section 4 estimates and tests heterogeneous risk aversion models and explores the ability of the models to explain the cross section of returns.

1.1 Related literature

For many years after Fisher, in analyzing the link between individual utility maximization and asset prices, financial theorists either employed a rational representative agent or considered cases of heterogeneous agent models that admit aggregation, that is, cases in which the model is equivalent to one with a representative agent. Extending the portfolio choice work of Markowitz (1952) and Tobin (1958), Sharpe (1964) and Lintner (1965a,b) established the Capital Asset Pricing Model (CAPM).[1] These original CAPM papers, which concluded that an asset's covariance with the aggregate market determines its return, actually allowed for substantial heterogeneity in endowments and risk preferences across investors. However, their form of quadratic or mean-variance preferences admitted aggregation and obviated the role of the wealth distribution.

The seminal consumption-based asset pricing work of Lucas (1978), Breeden (1979), and Hansen and Singleton (1983) also abstracted from investor heterogeneity. They and others derived and tested analytic relationships between the marginal rate of substitution of a representative agent (with standard preferences) and asset prices. Despite the elegance and tractability of the representative agent/aggregation approach, it has failed to adequately explain the fluctuations of asset prices in the economy. Largely inspired by the limited empirical fit of the CAPM (in explaining the cross section of stock returns), the equity premium puzzle (Mehra and Prescott, 1985), and excess stock market volatility and related price-dividend ratio anomalies (Shiller, 1981), since the 1980s theorists have extended macro/finance general equilibrium models to consider non-standard utility functions and meaningful investor heterogeneity. These models can be categorized into two groups.

In the first group, agents have identical standard (constant relative risk aversion) preferences but are subject to uninsured idiosyncratic risks.[2] Although the models of this literature have had some quantitative success, the empirical results (based on consumption panel data) are mixed and may even be spuriously caused by the heavy tails in the cross-sectional consumption distribution (Toda and Walsh, 2014).

In the second group, markets are complete and agents have either heteroge-

[1] See Geanakoplos and Shubik (1990) for a general and rigorous treatment of CAPM theory.

[2] Examples are Mankiw (1986), Constantinides and Duffie (1996), Heaton and Lucas (1996), Krusell and Smith (1998), Brav et al. (2002), Cogley (2002), Balduzzi and Yao (2007), Kocherlakota and Pistaferri (2009), among many others. See Ludvigson (2013) for a review.

neous CRRA preferences[3] or identical but non-homothetic preferences.[4] In this class of models the marginal rates of substitution are equalized across agents and a "representative agent" in the sense of Constantinides (1982) exists, but aggregation in the sense of Gorman (1953) fails. Therefore there is room for agent heterogeneity to matter for asset pricing. However, this type of agent heterogeneity is generally considered to be irrelevant for asset pricing because in dynamic models the economy is dominated by the richest agent (the agent with the largest expected wealth growth rate) in the long run (Sandroni, 2000; Blume and Easley, 2006).[5] One notable exception is Gârleanu and Panageas (2012), who study a continuous-time overlapping generations endowment economy with two agent types with Epstein-Zin constant elasticity of intertemporal substitution/constant relative risk aversion preferences. Even if the aggregate consumption growth is i.i.d. (geometric Brownian motion), the risk-free rate and the equity premium are time-varying, even in the long run. The intuition is that when the risk tolerant agents have a higher wealth share, they drive up asset prices and the interest rate. The effect of preference heterogeneity persists since new agents are constantly born. Consistent with our empirical findings and model, the calibration of Gârleanu and Panageas (2012) suggests that increasing the consumption share of more risk tolerant agents pushes down the equity premium. All of the above works are theoretical, and our paper seems to be the first in the literature to empirically test the asset pricing implications of models with preference heterogeneity.

Although the wealth distribution theoretically affects asset prices, there are few empirical papers that directly document this connection. To the best of our knowledge, Johnson (2012) is the only one that explores this issue using income/wealth distribution data. In particular, he shows that portfolios positively correlated with income concentration command lower average excess returns (consistent with our findings in Section 4). However, his analysis is quite different from ours. First, his model relies on a "keeping up with the Joneses"-type consumption externality with incomplete markets. In contrast, we employ a standard general equilibrium model (a plain vanilla Arrow-Debreu model). Next, he does not directly test moment conditions from his model, whereas we perform structural estimation of heterogeneous risk aversion parameters. Nor does he derive and estimate our two factor covariance pricing formula or discuss average risk tolerance. Finally, Johnson (2012) does not explore the ability of top income shares to predict market excess returns, and he detrends inequality differently from the way we do.

Lastly, our study is related to the empirical literature on heterogeneity in risk preferences. A number of recent papers have found that the wealthy have portfolios more heavily skewed towards risky assets, and many of these studies have concluded that the wealthy are relatively more risk tolerant, either due to declining relative risk aversion or innate heterogeneity in relative risk aversion. See, for example, Carroll (2002), Vissing-Jørgensen (2002), Campbell (2006), Bucciol and Miniaci (2011), or Calvet and Sodini (2014). This literature lends credibility to our premise that the rich are relatively more tolerant. Further-

[3]Examples are Dumas (1989), Wang (1996), Chan and Kogan (2002), Hara et al. (2007), Longstaff and Wang (2012), and Bhamra and Uppal (2014).

[4]Examples are Gollier (2001) and Hatchondo (2008).

[5]Coen-Pirani (2004) and Guvenen (2009) study cases with incomplete markets and heterogeneous Epstein-Zin preferences.

more, we structurally *estimate* that the rich are relatively more risk tolerant and thus support the findings of these authors, albeit with a different framework and dataset.

2 Asset pricing implications of preference heterogeneity

In this section we present two simple models in which the heterogeneity in agents' attitude towards risk matters for asset pricing and derive testable implications as well as a novel covariance pricing formula.

2.1 Asset pricing with heterogeneous risk aversion

2.1.1 Model with arbitrary preferences

Consider a two period model with time indexed by t and $t+1$. There are I agents indexed by $i = 1, \ldots, I$. Agent i has the expected utility over final wealth $w_{i,t+1}$,

$$\mathrm{E}_t[u_i(w_{i,t+1})],$$

where u_i is von Neumann-Morgenstern utility function with $u_i' > 0$ and $u_i'' < 0$. There are J assets indexed by $j = 1, \ldots, J$. Asset j trades at price q_j per share (to be determined in equilibrium) at t and pays dividend D_j at $t+1$. Agent i is endowed with n_{ij} shares of asset j at t. Let $w_{it} = \sum_{j=1}^{J} q_j n_{ij}$ be the initial wealth of agent i. Letting n_{ij}' be the number of shares agent i holds after trade, the optimal portfolio problem is

$$\begin{aligned} &\underset{\{n_{ij}'\}}{\text{maximize}} && \mathrm{E}_t[u_i(w_{i,t+1})] \\ &\text{subject to} && \sum_{j=1}^{J} q_j n_{ij}' = w_{it},\ w_{i,t+1} = \sum_{j=1}^{J} D_j n_{ij}'. \end{aligned} \tag{2.1}$$

Assuming no trade frictions, the first-order condition for optimality with respect to n_{ij}' is

$$\mathrm{E}_t[u_i'(w_{i,t+1})D_j] = \lambda_i q_j,$$

where $\lambda_i > 0$ is the Lagrange multiplier for the budget constraint. Dividing by q_j and letting $R_{j,t+1} = D_j/q_j$ be the gross return on asset j and assuming the existence of a risk-free asset (with gross risk-free rate $R_{f,t}$), we obtain

$$\mathrm{E}_t[u_i'(w_{i,t+1})(R_{j,t+1} - R_{f,t})] = 0.$$

Using the Taylor approximation

$$u_i'(x) \approx u_i'(a_i) + u_i''(a_i)(x - a_i)$$

around the expected future wealth $a_{it} = \mathrm{E}_t[w_{i,t+1}]$, letting $x = w_{i,t+1}$, we obtain

$$\mathrm{E}_t[(u_i'(a_{it}) + u_i''(a_{it})(w_{i,t+1} - a_{it}))(R_{j,t+1} - R_{f,t})] = 0, \tag{2.2}$$

where we have written $=$ instead of $\approx$. Dividing both sides by $-u_i''(a_{it}) > 0$ and using the definition of the relative risk tolerance (reciprocal of the Arrow-Pratt measure of relative risk aversion)

$$\tau_i = -\frac{u_i'(a_{it})}{a_{it}u_i''(a_{it})},$$

we obtain

$$\mathrm{E}_t[(a_{it}\tau_i - (w_{i,t+1} - a_{it}))(R_{j,t+1} - R_{f,t})] = 0. \tag{2.3}$$

Adding across all agents, letting $W_{t+1} = \sum_{i=1}^{I} w_{i,t+1}$ be the aggregate wealth at $t+1$, and dividing by $\mathrm{E}_t[W_{t+1}] = \sum_{i=1}^{I} a_{it}$, we obtain

$$\mathrm{E}_t[(\bar{\tau} - W_{t+1}/\,\mathrm{E}_t[W_{t+1}] + 1)(R_{j,t+1} - R_{f,t})] = 0,$$

where $\bar{\tau} = \sum_i a_{it}\tau_i / \sum_i a_{it}$ is the weighted average risk tolerance. Now since every asset must be held by some agent in equilibrium and there is no consumption at t, adding individual budget constraints, the growth rate of aggregate wealth must be equal to the *market return* $R_{m,t+1}$. Therefore $W_{t+1} = R_{m,t+1}W_t$. Taking expectations, we obtain $\mathrm{E}_t[W_{t+1}] = \mathrm{E}_t[R_{m,t+1}]W_t$. Therefore $W_{t+1}/\,\mathrm{E}_t[W_{t+1}] = R_{m,t+1}/\,\mathrm{E}_t[R_{m,t+1}]$. Putting all the pieces together, we obtain

$$\mathrm{E}_t[((\bar{\tau}+1)\,\mathrm{E}_t[R_{m,t+1}] - R_{m,t+1})(R_{j,t+1} - R_{f,t})] = 0, \tag{2.4}$$

which is the key moment condition that we will exploit throughout the rest of the paper.

Alternatively, if we apply the Taylor approximation around the initial wealth w_{it} instead of the expected future wealth $\mathrm{E}_t[w_{i,t+1}]$, (2.2) holds with $a_{it} = w_{it}$. Adding across i and dividing by aggregate wealth $W_t = \sum_i w_{it}$, we get

$$\mathrm{E}_t[(\bar{\tau} + 1 - R_{m,t+1})(R_{j,t+1} - R_{f,t})] = 0, \tag{2.5}$$

where $\bar{\tau} = \sum_i w_{it}\tau_i / \sum_i w_{it}$ is the average risk tolerance weighted by initial wealth.

The moment conditions (2.4) and (2.5) are both valid approximations; (2.4) is more accurate because we approximate around the expectation of the relevant variable (future wealth), but (2.5) is easier to handle because there is no need to predict future stock returns or the wealth distribution.

2.1.2 Model with quadratic preferences

The above discussion is only approximate since it involves linear approximations, but it can be made exact under some assumptions. Assume that in addition to the risky J assets, agents can trade a risk-free asset in zero net supply, where the risk-free rate R_f is determined in equilibrium. Let $\theta = (\theta_1, \ldots, \theta_J)$ be a portfolio, where θ_j is the fraction of wealth of a typical agent invested in asset j. Then the fraction of wealth invested in the risk-free asset is $1 - \sum_j \theta_j$. Therefore the return on the portfolio θ is

$$R(\theta) = \sum_{j=1}^{J} R_j\theta_j + R_f\left(1 - \sum_{j=1}^{J}\theta_j\right) = R_f + \sum_{j=1}^{J}(R_j - R_f)\theta_j,$$

where R_j is the gross return on asset j.

Now consider exactly the same model except preferences. Suppose that agents are mean-variance optimizers. More precisely, agent i maximizes

$$v_i(\theta) = \mathrm{E}[R(\theta)] - \frac{1}{2\tau_i}\,\mathrm{Var}[R(\theta)],$$

where $\tau_i > 0$ is the risk tolerance. The expected return and variance of the portfolio are

$$\mathrm{E}[R(\theta)] = R_f + \langle \mu - R_f 1, \theta \rangle\,, \;\; \mathrm{Var}[R(\theta)] = \langle \theta, \Sigma\theta \rangle\,,$$

respectively, where $\mu = (\mu_1, \ldots, \mu_J)$ is the J-vector of expected returns $\mu_j = \mathrm{E}[R_j]$, 1 is the J-vector of ones, and Σ is the variance-covariance matrix of the returns $\mathbf{R} = (R_1, \ldots, R_J)$. Thus the optimal portfolio problem of agent i reduces to

$$\underset{\theta}{\text{maximize}} \;\; R_f + \langle \mu - R_f 1, \theta \rangle - \frac{1}{2\tau_i}\,\langle \theta, \Sigma\theta \rangle\,,$$

where $\theta \in \mathbb{R}^J$ is unconstrained. The first-order condition is

$$\mu - R_f 1 - \frac{1}{\tau_i}\Sigma\theta = 0 \iff \theta_i^* = \tau_i \Sigma^{-1}(\mu - R_f 1), \tag{2.6}$$

where θ_i^* is the optimal portfolio of agent i.

Since every asset must be held by someone and the risk-free asset is in zero net supply by definition, the average portfolio weighted by individual wealth, $\sum w_{i0}\theta_i^* / \sum_i w_{i0}$, must be the *market portfolio* (value-weighted average portfolio), denoted by θ_m. Letting $\bar\tau = \sum_{i=1}^I w_{i0}\tau_i / \sum_{i=1}^I w_{i0}$ be the average risk tolerance, by taking the weighted average of the first-order condition (2.6), we get

$$\theta_m = \bar\tau \Sigma^{-1}(\mu - R_f 1). \tag{2.7}$$

Multiplying both sides of (2.7) by Σ and comparing the j-th element, we get

$$\mathrm{Cov}[R_m, R_j] = \bar\tau\,\mathrm{E}[R_j - R_f],$$

where $R_m = R(\theta_m)$ is the market return. Since

$$\begin{aligned}\mathrm{Cov}[R_m, R_j] &= \mathrm{E}[(R_m - \mathrm{E}[R_m])(R_j - R_f + R_f)] \\ &= \mathrm{E}[(R_m - \mathrm{E}[R_m])(R_j - R_f)],\end{aligned}$$

it follows that

$$\mathrm{E}[(\bar\tau + \mathrm{E}[R_m] - R_m)(R_j - R_f)] = 0. \tag{2.8}$$

(2.8) is identical to (2.4) or (2.5) except that 1 is replaced by $\mathrm{E}[R_m]$.

2.2 Covariance pricing formula

Using the moment conditions derived above, we can obtain a two factor covariance pricing formula. For simplicity, assume that there are two types of agents with high and low risk tolerance τ_H and τ_L. Letting

$$0 < \alpha_t = \mathrm{E}_t[w_{H,t+1}] / \mathrm{E}_t[W_{t+1}] < 1$$

be the fraction of wealth of high risk tolerant agents, from (2.4) we get

$$\mathrm{E}_t[((1+\alpha_t\tau_H+(1-\alpha_t)\tau_L)\,\mathrm{E}_t[R_{m,t+1}]-R_{m,t+1})(R_{j,t+1}-R_{f,t})]=0. \quad (2.9)$$

We can derive similar moment conditions from (2.5) and (2.8).

The moment condition (2.9) implies that

$$M_{t+1}=(1+\alpha_t\tau_H+(1-\alpha_t)\tau_L)\,\mathrm{E}_t[R_{m,t+1}]-R_{m,t+1}$$

is a scaled stochastic discount factor. Taking the conditional expectations and using the definition of covariance, from (2.9) with the market return we obtain

$$\mathrm{E}_t[R_{m,t+1}]-R_{f,t}=-\frac{\mathrm{Cov}_t[M_{t+1},R_{m,t+1}]}{\mathrm{E}_t[M_{t+1}]}=\frac{1}{\tau_L+\alpha_t\Delta\tau}\frac{\mathrm{Var}_t[R_{m,t+1}]}{\mathrm{E}_t[R_{m,t+1}]}, \quad (2.10)$$

where $\Delta\tau=\tau_H-\tau_L>0$. Similarly, taking the unconditional expectations and using the definition of covariance, from (2.9) we obtain

$$\begin{aligned}\mathrm{E}[R_{j,t+1}]-\mathrm{E}[R_{f,t}]&=-\frac{\mathrm{Cov}[M_{t+1},R_{j,t+1}]}{\mathrm{E}[M_{t+1}]}\\&=\frac{1}{\mathrm{E}[M_{t+1}]}\,\mathrm{Cov}[R_{m,t+1},R_{j,t+1}]-\frac{\Delta\tau}{\tau_L+\mathrm{E}[\alpha_t]\Delta\tau}\,\mathrm{Cov}[\alpha_t,R_{j,t+1}]. \quad (2.11)\end{aligned}$$

By (2.10) and (2.11), we obtain the following proposition.

Proposition 2.1. *Suppose that the rich are more risk tolerant than the poor. Then a high top wealth share predicts a low equity premium. Furthermore, the top wealth share is an asset pricing factor.*

The covariance pricing formula (2.11) shows that the covariance both with the market *and* with the wealth distribution are priced. In particular, assets with returns that are positively correlated with the wealth share of more risk tolerant agents should have lower average returns. As we outlined in Section 1.1, a number of recent empirical studies show that the rich choose riskier portfolios and argue that the wealthy have lower relative risk aversion than do poorer investors. We estimate the same pattern in Section 4. To the extent that the richest agents are the most risk tolerant, assets that are positively correlated with the wealth share of the rich should have lower average returns. In short, the theory produces a two factor model of asset prices.

Why does negative correlation with the wealth share of risk tolerant agents lead to high average returns for an asset? Intuitively, an asset that pays off when the wealth share of risk tolerant agents is low is delivering precisely when the economy's average marginal utility is low. This is because with, for example, mean-variance or CRRA preferences, marginal utility is increasing in risk tolerance, all else equal. Thus, when risk tolerance is low on average, marginal utility is low. Moreover, assets that deliver in low marginal utility times will on average be less attractive to the agents. Therefore, the agents demand on average a higher risk premium on assets that pay off when the risk tolerant wealth share is low. In short, the implicit representative or average agent has high marginal utility under two scenarios: low market returns and high risk tolerance. Assets that deliver in these states provide insurance to this representative agent and thus command lower risk premiums on average.

The covariance pricing formula (2.11) still contains the expected future wealth share $\alpha_t = \mathrm{E}_t[w_{H,t+1}]/\mathrm{E}_t[W_{t+1}]$ and therefore is not directly amenable to data. However, as long as the top wealth share is slowly moving relative to asset returns (which is the case), we can replace it by the *actual* wealth share $w_{H,t+1}/W_{t+1}$. In fact, this approximation holds *exactly* in a continuous-time model studied by Gârleanu and Panageas (2012).

2.3 Numerical example

In this section we numerically solve two examples of the model in Section 2.1, one with agents with constant but heterogeneous relative risk aversion (CRRA) and another with identical decreasing relative risk aversion (DRRA) agents except initial wealth. Appendix B discusses the numerical algorithm in detail.

2.3.1 Two CRRA agents with heterogeneous risk aversion

Assume that there are two agent types, $i = 1, 2$. Agent 1 has high risk tolerance τ_H and agent 2 has low risk tolerance τ_L. For numerical values, we set $\gamma_H = 1/\tau_H = 0.5$ and $\gamma_L = 1/\tau_L = 2$. There is only one risky asset (stock) and a risk-free asset in zero net supply. Fraction α of stocks are initially held by agent 1 and fraction $1 - \alpha$ by agent 2. There are two states with equal probability, and the dividend of the stock is $1 + \mu \pm \sigma$, where $\mu = 0.07$ and $\sigma = 0.2$. To see the accuracy of the approximation, we both solve the exact model numerically as well as the approximate model semi-analytically using either (2.4) or (2.5).

The results are shown in Figure 1. `A1` and `A2` refer to the approximate model using (2.4) and (2.5), respectively. According to Figure 1a, the optimal portfolio of the exact and the two approximate models are close, at least when the wealth share of the risk tolerant agent 1 is not too small. As the risk tolerant agent gets richer, the risk averse agent's portfolio share of stock declines. Essentially agent 1 is providing insurance to agent 2.

According to Figure 1b, contrary to the case with portfolios the equilibrium equity premium of `A2` is not so accurate. The approximation error can be up to 2% in magnitude. However, the approximation `A1` is virtually indistinguishable from the exact model. As the risk tolerant agent gets richer, there is more demand for borrowing, and therefore the risk-free rate increases in order to clear the market. In this example since the expected stock return is fixed at 7%, the equity premium shrinks as the risk tolerant agent gets richer.

2.3.2 Two agents with identical DRRA utilities

Consider the same example as above except that preferences are identical and exhibit decreasing relative risk aversion (DRRA). It is natural to assume that the Arrow-Pratt measure of relative risk aversion is a decreasing power function,

$$\mathrm{RRA}(x) = -\frac{xu''(x)}{u'(x)} = \gamma \left(\frac{x}{c}\right)^{-\eta},$$

where $\gamma, \eta, c > 0$ are parameters.[6] The economic interpretation of the parameters is that c is a reference point for wealth, γ is the relative risk aversion

[6] This specification has essentially two parameters, since only η and γc^{η} are identified. $\eta = 0$ corresponds to the CRRA case.

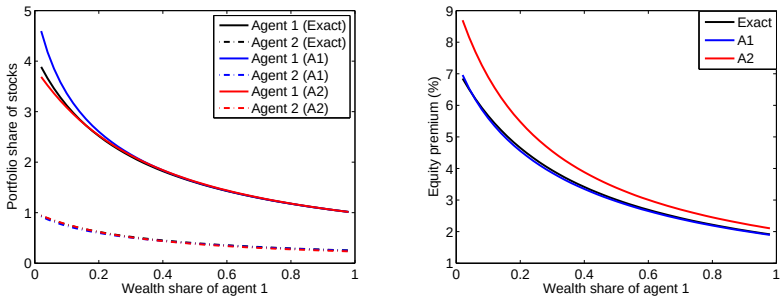

(a) Portfolio share of stock. (b) Equity premium.

Figure 1: Numerical example with heterogeneous CRRA preferences. `Exact`: numerical solution of exact model; `A1`, `A2`: semi-analytical solution of approximate model using (2.4) and (2.5).

coefficient at this reference point, and η governs the speed (elasticity) at which RRA decreases. Solving the ordinary differential equation, it follows that the von Neumann-Morgenstern utility function is

$$u(x) = A \int_c^x \mathrm{e}^{\frac{\gamma}{\eta}\left(\frac{y}{c}\right)^{-\eta}} \mathrm{d}y + B,$$

where $A > 0$ and B are some constants. Since A and B merely define an affine transformation, they do not affect agents' behavior. Therefore, without loss of generality we may assume $A = 1$ and $B = 0$, so the utility function is

$$u(x) = \int_c^x \mathrm{e}^{\frac{\gamma}{\eta}\left(\frac{y}{c}\right)^{-\eta}} \mathrm{d}y. \tag{2.12}$$

For a numerical example, we normalize the aggregate wealth at $t = 0$ to be $W_0 = 1$ and set $\gamma = 2$, $\eta = 1$, and $c = 1/2$ (the reference point is equal distribution of wealth), so $u(x) = \int_{1/2}^x \mathrm{e}^{1/y} \mathrm{d}y$ and $\mathrm{RRA}(x) = 1/x$. Figure 2 shows the numerical solution. According to Figure 2a, as agent 1 gets richer, he becomes less risk averse and invests more in stocks. However, when he is too rich agent 2 is too poor to lend, and agent 1's portfolio share of stocks eventually decreases. Consistent with empirical evidence discussed in Section 1.1, the wealthy choose risker portfolios. According to Figure 2b, the equity premium is highest when the wealth is equally distributed. As the wealth distribution becomes more skewed, the richer and more risk tolerant agent leverages and drives down the equity premium. As in the previous example, the approximation `A1` is excellent but `A2` is poor.

3 Empirical link between inequality and equity premium

Thus far, we have theoretically analyzed some models and examples in which the extent of inequality across agents with heterogeneous risk aversion is key in predicting returns. We found not only that the wealth distribution affects

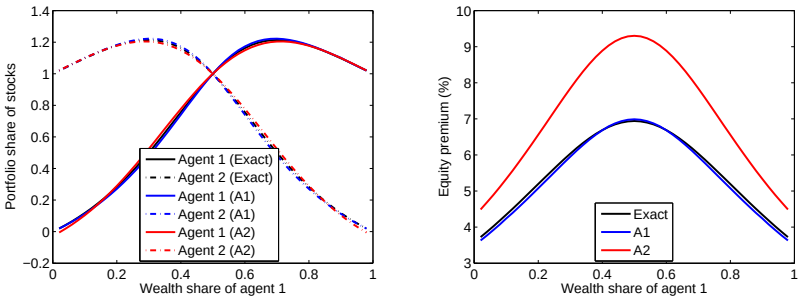

(a) Portfolio share of stock.

(b) Equity premium.

Figure 2: Numerical example with identical DRRA preference. `Exact`: numerical solution of exact model; `A1`, `A2`: semi-analytical solution of approximate model using (2.4) and (2.5).

the relative prices of risky assets but also that the extent of inequality may determine an economy's overall risk premium (and thus the equity premium).

But, are macroeconomic and financial data consistent with the implications of this paper's model and those in the above literature? In this section, we show that there is a strong and robust negative relationship between the top income/wealth share and subsequent medium-term excess stock market returns. That is, current inequality appears to forecast the risk premium of the U.S. stock market. The negative sign of the relationship is consistent with the above example, and the inequality measures do not seem to merely be proxying for either of two leading predictors of excess returns, price-dividend ratio (Shiller, 1981) and the consumption-wealth ratio (Lettau and Ludvigson, 2001).

3.1 Data

We employ the Piketty and Saez (2003) inequality measures for the U.S., which are available in spreadsheets on the website of Emmanuel Saez.[7] In particular, we consider top income and wealth share measures. The income measures (with or without capital gains) are at the annual frequency and are based on tax return data, and cover the period 1913–2012. The wealth series, the top 1% wealth share, covers 1916–2000 at the annual frequency and is based on estate tax data. As opposed to the income data, many years are missing in the 50s, 60s, and 70s for the wealth data, so we complete the series with cubic interpolation. The income series reflect in a given year the percent of income earned by the top earners pretax. Similarly, the wealth series is the percent of wealth owned by the richest 1%. See Piketty and Saez (2003) and Kopczuk and Saez (2004) for further details on the construction of these series. In Appendix A, we also consider .1% shares and the wealth share series of Saez and Zucman (2014).

Figure 3a shows the top income share for each group, the richest 0–0.5%, 0.5–1%, 1–5%, and 5–10%. All groups seem to share a common trend, which is similar to the highest marginal tax rate in Figure 4. However, the behavior of these series around the trend is quite different. First, the top 0.5–1% share is very smooth. Second, the top 0–0.5% share seems procyclical (move in the

[7] `http://elsa.berkeley.edu/~saez/`

same direction as booms and busts), which is most apparent in the 1920s, 1960s, 1990s, and mid-2000s. On the other hand, the behavior of the top 1–5% and 5–10% resemble each other and seems countercyclical (move in the *opposite* direction as booms and busts). Figure 3b shows the relative income share of each group within the top 10%. We can see that the top 0.5–1% is stable, the top 0–0.5% moves in the same direction as booms and busts, and the top 1–10% moves in the opposite direction.

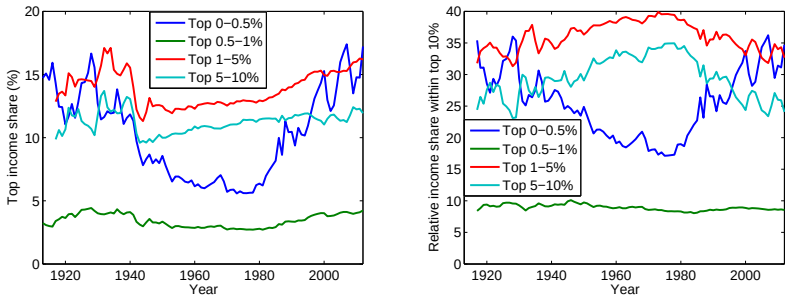

(a) Absolute shares. (b) Relative shares within the top 10%.

Figure 3: Top income shares including capital gains (1913–2012).

Within the context of the model in Section 2, this behavior can be explained if the richer agents are more risk tolerant. Consider, for example, the mean-variance model. Then the mutual fund theorem holds and agents invest more or less than 100% in stocks according as whether they are more or less risk tolerant than the average. If we assume that the top 10% hold the entire stock market, Figure 3b tells us that the top 0.5–1% roughly hold the market portfolio, the top 0–0.5% are more risk tolerant and leverage (borrow from the poor), and the top 1–10% are more risk averse and lend to the richest 0.5%. Thus, in bringing our theory to the data, we take the top 1%, rather than say the top 5% or 10%, as our dividing line between more and less risk tolerant agents.

Below, in both regressions and GMM exercises, we use not the raw Piketty-Saez series but rather detrended, stationary versions. Specifically, we detrend each of the inequality measures using the Hodrick-Prescott (HP) filter with a smoothing parameter of 100, which is standard for annual frequencies. Stochastically detrending asset return predictors is in the tradition of Campbell (1991), for example, who removes a trend in the short-term interest rate before including it in stock return vector autoregressions. Indeed, the Piketty-Saez series appear to exhibit a U-shaped trend over the century, which might be due to the change in the marginal income tax rates. According to Figure 4, the marginal tax rate for the highest income earners increased from about 25% to 90% over the period 1930–1945 and started to decline in the 1960s, reaching about 40% in the 1980s.[8] Thus the marginal tax rate exhibits an inverse U-shape that coincides with the trend in the Piketty-Saez series. Imposing stationarity in this way helps ensure the validity of standard error calculations and inference and prevent spurious regressions. Figure 5 plots the top 1% series (with capital gains) and their estimated trends. In Appendix A we employ different filtering methods. Our results appear robust to using a smoothing parameter of 10,

[8]The tax rate data is from the Tax Foundation (`http://taxfoundation.org/`).

the one-sided HP filter, the Kalman filter (which is also one-sided), the moving average filter, or linear detrending.

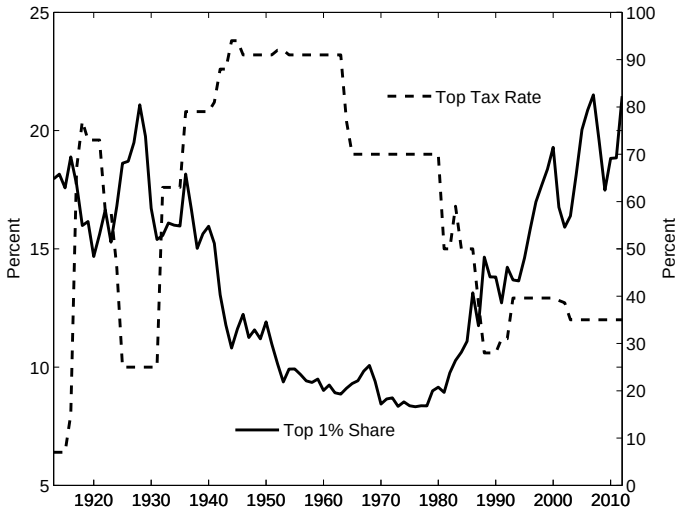

Figure 4: Top 1% income share including capital gains (left axis) and top marginal tax rate (right axis), 1913–2012.

We calculate annual one and five year U.S. stock market excess returns using the annual stock market spreadsheet from the website of Robert Shiller.[9] The spreadsheet contains historical one year interest rates and price, dividend, and earnings series for the S&P 500 index, which are all put into real terms using the consumer price index (CPI). These data are also used to calculate the series P/E10 and P/D10, which are the price-dividend and price-earnings ratios (in real terms) for the S&P 500 based on 10 year moving averages of earnings and dividends.

For the Lettau-Ludvigson consumption-wealth ratio, commonly referred to as CAY, we use 100 times the annual version of this series from the website of Amit Goyal (the spreadsheet for Welch and Goyal (2008)).[10] It spans the period 1945–2012.

We also include as controls GDP growth and, inspired by Lettau et al. (2008) and Bansal et al. (2014), consumption volatility. Annual data for real GDP and real consumption are from the website of the Federal Reserve Bank of St. Louis[11] and span 1930–2012. We estimate consumption volatility using a standard GARCH(1,1) model for consumption growth.

Finally, we investigate the relationship between the top 1% income share and the three annual Fama-French asset pricing factors from the website of Kenneth French: EP, SMB, and HML (1927–2012).[12] EP is the value-weighted excess return on U.S. CRSP firms. SMB is, roughly, the difference in return between

[9] `http://www.econ.yale.edu/~shiller/data.htm`
[10] `http://www.hec.unil.ch/agoyal/`
[11] `http://research.stlouisfed.org/fred2/`
[12] `http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/index.html`

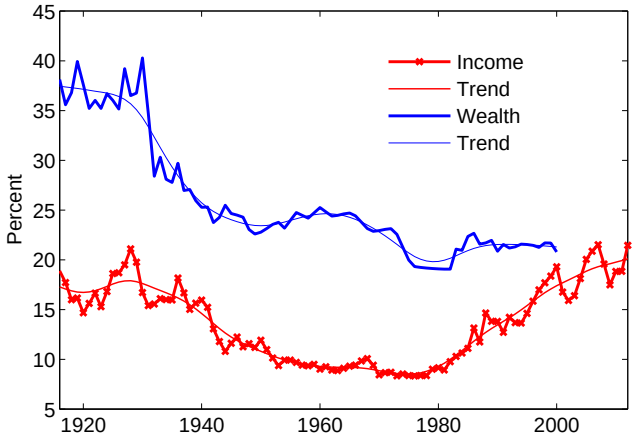

Figure 5: Top 1% income share including capital gains (1913–2012) and top 1% wealth share (1916–2000). The thin lines are the HP filter trends.

small and large valued firms. HML is, roughly, the difference in return between high and low book-to-market firms. See also Fama and French (1993).

3.2 Regression analysis

Table 1 shows the results of regressions of one year (t to $t+1$) excess stock market returns on top share measures (time t) and some classic return predictors (time t).[13] We find that when the top 1% income share (January to December of year t) rises above trend by one percentage point, subsequent one year market excess returns (January to December of year $t+1$) decline on average by 5.6%. The coefficient is significant at the 1% level (using Newey-West standard errors), and the R-squared statistic is .09. Figure 6 shows the corresponding scatter plot for five year returns. It is clear, at least in sample, that the detrended top 1% share series has substantial power in forecasting the subsequent overall excess return on the stock market.

This relationship also holds for the top 1% wealth share. With respect to one year returns, the top wealth share is strongly significant and yields an R-squared of .24, which is greater than the income share R-squared.

Given the strength of the relationship, a question immediately arises. Is there some mechanical, non-equilibrium explanation for the relationship between inequality and subsequent excess returns? For example, might stock returns somehow be *determining* the top share measures? For a few reasons, the answer is likely no. First, the relationship is between *initial* inequality and *subsequent* returns. Returns could affect contemporaneous top shares but likely not lagged top shares. One might still worry that our results are driven by the HP filter, which uses the past, current, and future data to obtain a smooth trend, thereby introducing a look-ahead bias.[14] As mentioned above, in Appendix A

[13] In Table 7 in Appendix A we use five year (t to $t+5$) returns.

[14] For example, since the rich are likely to be more exposed to the stock market, when the stock market goes up at year $t+1$, the rich will be richer than usual. But then the trend in

we also detrend the top income share by the Kalman filter, the one-sided HP filter, and the moving average filter, which use only past information and obtain similar results. Finally, as we see in regression (2) from Tables 1 and 7, when excluding capital gains, the top 1% income share coefficient is *larger* with respect to five year returns and only slightly smaller with one year returns. If returns were strongly affecting lagged inequality, excluding capital gains would likely mitigate the regression results.

But, one might say, we have known at least since Shiller (1981) that when prices are high relative to either earnings or dividends, subsequent market excess returns are low. The current price could indeed affect current inequality (see Section 3.3). Are the top shares series simply proxying for the price-dividend or price-earnings ratios, which are known to predict returns? Again, the answer seems to be no for two reasons. First, excluding capital gains from income does not significantly mitigate the relationship, and capital gains are the main avenue through which prices would determine inequality. Second, as we see in regressions (6) and (7) from Table 1, top shares predict excess returns *even when controlling for the log price-dividend or price-earnings ratio*. Including these controls does decrease the top shares coefficients slightly, but they remain large and significant. In the case of one year returns, the P/D and P/E ratios are not significant after controlling for top income shares.

In regressions (4), (5), (8), and (9) from Table 1, we also control for real GDP growth, consumption volatility (Lettau et al. (2008) and Bansal et al. (2014)), and CAY, which Lettau and Ludvigson (2001) show forecasts market excess returns. In the case of one year returns, including these controls has little impact on the relationship between the top income share and subsequent returns.

Our empirical analysis thus far has relied on detrending, which requires the researcher to take a stand on smoothing parameters and the underlying trend model. Do the raw data indicate a relationship between asset prices and the one percent? Figure 7 suggests that the answer is the yes. Over 1913-2012, both overall and within subsamples, there is a clear positive correlation between the top 1% income share (not detrended) and the price-dividend ratio. Of course, this scatter plot does not establish causation, but it is more evidence in favor of our theory and suggests that our empirical results are not simply artifacts of detrending. Indeed, as we have shown, above trend inequality predicts subsequent excess returns even when using a simple, one-sided trend estimation method like the ten year moving average.

In summary, the data appear consistent with our theory that an increasing concentration of income or wealth decreases the market risk premium.

the top income share will shift upwards, and the year t deviation of the top income share will be lower. Therefore the low income share at year t may spuriously predict a high stock return at $t+1$.

Table 1: Regressions of one year excess stock market returns on top income shares and other predictors

	Dependent Variable: t to $t+1$ Excess Stock Market Return								
Regressors (t)	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Constant	6.66	6.69	7.09	3.90	8.41	21.60	21.83	6.28	32.28
	(1.77)	(1.76)	(1.83)	(2.36)	(3.03)	(11.65)	(12.86)	(1.64)	(19.96)
Top 1%	-5.61***			-7.63***	-7.67***	-5.02***	-4.86***	-5.50***	-5.53***
	(1.60)			(1.82)	(1.68)	(1.63)	(1.71)	(1.56)	(1.61)
Top 1% (no cg)		-5.16**							
		(2.25)							
Top 1% (wealth)			-7.93***						
			(1.91)						
Real GDP Growth				0.73*					0.34
				(0.37)					(0.59)
Cons. Growth Volatility					-0.46				-1.46
					(0.69)				(2.34)
log(P/D10)						-4.52			-6.91
						(3.40)			(4.89)
log(P/E10)							-5.62		
							(4.66)		
CAY								2.39***	2.34***
								(0.58)	(0.68)
Sample	1913-	1913-	1916-	1930-	1931-	1913-	1913-	1945-	1945-
	-2012	-2012	-2000	-2012	-2012	-2012	-2012	-2012	-2012
R^2	.09	.05	.24	.14	.15	.10	.11	.22	.24

Newey-West standard errors in parentheses ($k = 4$)
Consumption growth volatility from GARCH(1,1) model
***, **, and * indicate significance at 1%, 5%, and 10% levels (suppressed for constants)

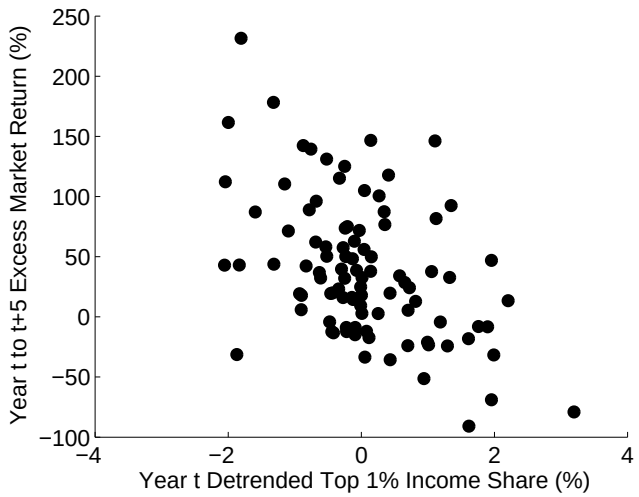

Figure 6: Year t to year $t+5$ excess stock market return vs. year t detrended top 1% income share including capital gains (1913–2008).

3.3 Relationship with return predictors and taxes

As we saw in Section 3.2, controlling for the price-dividend (or price-earnings ratio) or CAY mitigates to a small degree the estimated effect of inequality on subsequent excess returns. Furthermore, because the rich hold more stock than do the poor, high prices and the resulting capital gains likely have some direct impact on the top income shares. To what extent then are the top income shares correlated with classic return predictors? In Table 2, we regress the top 1% share on a number of series known to predict or explain asset returns (and on the top tax rate).

For the top 1% share, the correlation with the log price-dividend ratio is significant, but the R-squared is only .07. Therefore, while correlation with the price-dividend ratio likely explains some of the relationship between inequality and subsequent returns, it is not all or even most of the story. CAY, however, is not significantly correlated with the top 1% share. Overall, the top 1% income share appears to represent a component of the equity premium orthogonal to CAY and only slightly related to the price-dividend ratio.

Unsurprisingly given Figure 4, the 1% share has a strongly significant negative relationship with the detrended top marginal tax rate, and the corresponding R-squared is .10.

Table 2 also displays regressions of the top income share on the Fama-French factors, which are powerful in explaining the cross section of average returns. Notably, the SMB coefficient is significant, and the R-squared is larger than in the price-dividend and CAY regressions.

This result opens up the possibility that Fama-French factors—which are empirically motivated and theoretically hard to interpret—might actually be capturing the income/wealth inequality and hence the average risk aversion in the economy. The implication is that, across assets, heterogeneous correlation with top income shares may help explain heterogeneity in average returns. We

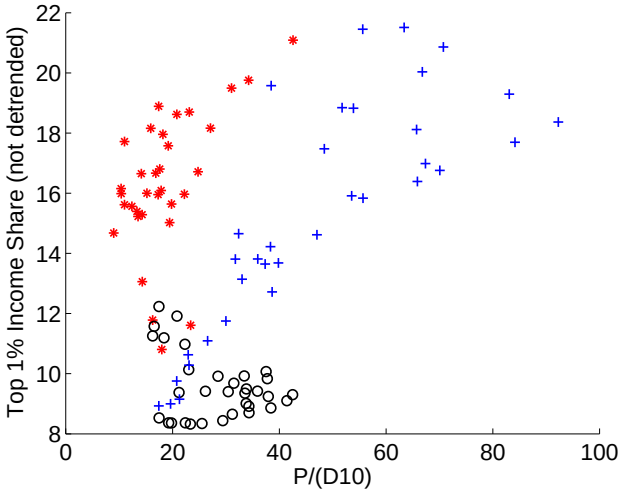

Figure 7: Top 1% income share (not detrended) vs. price-dividend ratio (in real terms) for the S&P 500 based on 10 year moving averages of dividends. 1913-1945 (*), 1946-1978 (o), and 1979-2012 (+).

Table 2: Regressions of top income share on factors

	Dependent Variable: Top 1% Share			
Regressors	(1)	(2)	(3)	(4)
Constant	-1.74	-1.07	-0.04	0.07
	(0.82)	(0.93)	(0.013)	(0.012)
log(P/D10)	0.53**	0.31		
	(0.25)	(0.28)		
CAY		-0.02		
		(0.05)		
Top Tax Rate♣			-0.04***	
			(0.01)	
EP				0.01
				(0.01)
SMB				-0.03**
				(0.01)
HML				-0.01
				(0.01)
Sample	1913-	1945-	1913-	1927-
	-2012	-2012	-2012	-2012
R^2	.07	.03	.10	.13

Newey-West standard errors in parentheses ($k = 4$)
♣ Detrended with HP filter ($\lambda = 100$)
***, **, and * indicate significance at 1%, 5%, and 10% levels (suppressed for constants)

explore this point in Section 4.

3.4 Instrumental variables regressions

As we saw in Table 2, the 1% income share is significantly correlated with the detrended top marginal tax rate (with an R-squared of .10). The tax series is, however, not significantly correlated with subsequent excess returns. Therefore, as an additional robustness check, we use the top tax rate as an instrument for the 1% share (Table 3). For one year returns (column (2)), including the instrument magnifies the effect of inequality on subsequent asset returns from -5.61% per year to -6.44% but also blows up the standard error, leading to statistical insignificance. For five year returns (column (5)), the inequality coefficient similarly becomes more negative but remains significant at the 1% level.

Given the relationship between log(P/D10) and the 1%, in columns (3) and (6) we also include the price-dividend ratio as an instrument. With this specification, inequality has a larger impact on subsequent return and is significant at the 10% level both with respect to one and five year returns.

Table 3: Instrumental Variables Regressions

	Dependent Variable: Excess Stock Market Return					
	t to $t+1$			t to $t+5$		
	(1)	(2)	(3)	(4)	(5)	(6)
Regressors (t)						
Constant	6.66	6.65	6.08	38.44	38.27	38.40
	(1.82)	(1.84)	(1.78)	(5.09)	(5.27)	(5.65)
Top 1%	-5.61***	-6.44	-10.49*	-25.12***	-37.08***	-48.17***
	(1.94)	(6.91)	(5.45)	(5.71)	(12.75)	(12.53)
Instruments (t)						
Top Tax Rate♣		✓	✓		✓	✓
log(P/D10)			✓			✓
R^2	.09	.09	.02	.20	.16	.03

Sample: 1913-2012
Two-step GMM standard errors in parentheses
♣ Detrended with HP filter ($\lambda = 100$)
***, **, and * indicate significance at 1%, 5%, and 10% levels
(suppressed for constants)

4 Testing the asset pricing implications for the cross section of returns

In this section we estimate and test a simplified version of the asset pricing implications derived in Section 2, namely the moment conditions (2.4), (2.5), and (2.8) with two agent types with risk tolerance $\tau_H > \tau_L$. For future reference, we refer to these models as follows.

`A1` The moment condition (2.9) derived from (2.4), which comes from the CRRA model with Taylor approximation at future expected wealth.

`A2` The moment condition derived from (2.5), which comes from the CRRA model with Taylor approximation at current wealth.

`MV` The moment condition (2.8), which comes from the mean-variance model.

We perform the structural GMM estimation as well as the Fama and MacBeth (1973) two-pass regression.

4.1 Data

Since the Piketty-Saez data is annual, we use annual asset returns data from 1927 to 2012. R_m is the CRSP value-weighted average portfolio return. R_j's are the 6 Fama-French portfolios sorted by size and book-to-market ratio and the 5 industry portfolios. R_f is the annualized return of the 90 day T-Bill rate.[15] Nominal returns are converted to real returns using the CPI. Asset returns and inflation data are from CRSP and Kenneth French's website.

Let top1_t be the detrended top 1% income share with capital gains at year t. Therefore top1_t has roughly mean zero and moves around zero, so τ_L in (2.9) can be interpreted as the 'baseline' risk tolerance in the economy and $\Delta\tau = \tau_H - \tau_L$ can be interpreted as the sensitivity of the risk tolerance on the top income share. We use the top 1% income share data, not the wealth share data, which is the theoretically relevant variable. We have two justifications for this choice. First, most income of the very rich is capital income, which should be proportional to wealth. Second, the raw income share data is available annually, unlike the wealth data which has either many missing years or is imputed.

4.2 GMM

4.2.1 Estimation

We estimate each model by GMM (with the market return and the Fama-French portfolios as test assets) using the identity matrix for the first stage estimation and the Newey-West HAC estimator with 4 lags to compute standard errors.[16] In estimating Models `A1` or `MV`, we need to compute the expected market return $\text{E}[R_m]$. For this purpose, inspired by regression (1) in Table 1 we regress R_m (year t to $t+1$) on a constant and top income share (year t) and define $\text{E}[R_m]$ to be the OLS fitted value. Following the discussion in Section 2.1, we adopt the following time convention. In the case of Model `A1`, for instance, we estimate

$$\text{E}[((1 + \tau_L + (\tau_H - \tau_L)\text{top1}_t)\,\text{E}[R_{m,t}] - R_{m,t})(R_{j,t} - R_{f,t})] = 0.$$

That is, we employ the detrended Piketty-Saez inequality measure top1_t and asset returns of the same year.

[15] Strictly speaking, we should use the 1 year bond return, but it is available only after 1941. We also estimate our model using the imputed value of the 1 year bond return before 1941 by regressing the 1 year bond return on a constant, 90 day T-bill rate, 30 day T-bill rate, and inflation, but the results were almost identical.

[16] We do not use the efficient second stage estimation results since it is well-known that the finite sample property is poor, see Cochrane (2005).

4.2.2 Results

Table 4 shows the results of the first stage GMM estimation of (2.9). The results are roughly the same across specifications. The estimated risk tolerance shows that the rich agents are nearly risk-neutral ($\gamma_H = 1/\tau_H \approx 0$) and the poor agents have a relative risk aversion coefficient $\gamma_L = 1/\tau_L$ in the range 2 to 3. The 'rich' and 'poor' here actually refer to the 'very rich' and 'ordinary stock market participant' in common language usage. Therefore our 'poor' agents are still relatively rich compared to the whole population. In any case, these results validate our claim that the 1% are more risk tolerant.

According to the J test (shown by P_J in Table 4), the models cannot be rejected, although this is not surprising given the small sample size ($T = 86$). Although the point estimates of the risk tolerance of each type are quite different, since the standard error of the risk tolerance of the rich is large, the Wald test (shown by $P(\tau_H = \tau_L)$ in Table 4) fails to reject homogeneous risk aversion.

Table 4: 1st stage GMM estimation of the moment conditions. Newey-West standard errors in parentheses ($k = 4$). Sample: 1927–2012.

Income share	Top 1%		
Model	`A1`	`A2`	`MV`
τ_H (rich)	17.9	22.3	20.3
	(14.4)	(17.3)	(15.5)
τ_L (poor)	0.378	0.556	0.423
	(0.090)	(0.124)	(0.106)
$P(\tau_H = \tau_L)$	0.22	0.21	0.20
P_J	0.24	0.32	0.27

Figure 8 shows the scatter plot of predicted and realized excess returns of each Fama-French portfolio as well as the market portfolio for each model. The scatter plot lies almost on the 45 degree line for models with heterogeneous risk aversion (Figures 8a–8c). However, when we estimate the standard CAPM by imposing $\tau_H = \tau_L$, the model fits poorly (Figure 8d). These results suggest that preference heterogeneity indeed matters for asset pricing.

Figure 9 shows the time series of the implied average risk tolerance

$$\bar{\tau}_t = \text{top1}_t \tau_H + (1 - \text{top1}_t)\tau_L$$

for each model. We can see that the average risk tolerance of the economy is generally around 0.5 but rises in booms (1920s, 1960s, 1990s, around 2005) and approaches (sometimes hits) zero in busts, consistent with Fisher's story that booms and busts are associated with the wealth distribution and therefore the average risk tolerance of the economy.

The Fisher story is most visible by plotting the average risk aversion. Figure 10 shows the time series of the implied average risk aversion $\bar{\gamma} = 1/\bar{\tau}$. We can see that the average risk aversion is usually between 2 and 4, but it sharply rises during bad times for the rich (the financial crises of early 1930 and 2007–2009, the introduction of exorbitant income tax during World War II, the collapse of the IT bubble in 2000–2002, and, to some extent, the income tax hike in 1993).

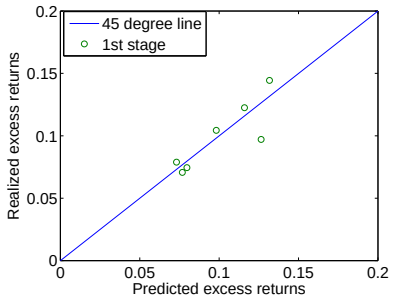

(a) Model `A1`

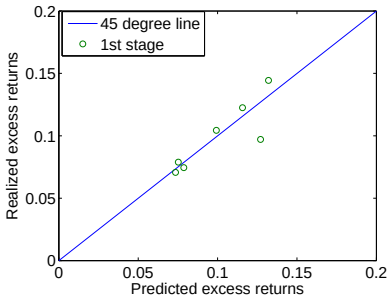

(b) Model `A2`

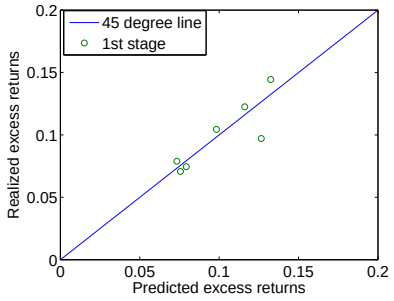

(c) Model `MV`

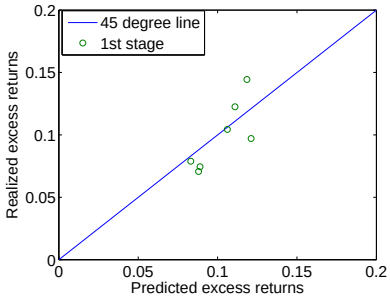

(d) Standard CAPM.

Figure 8: Predicted and realized excess returns.

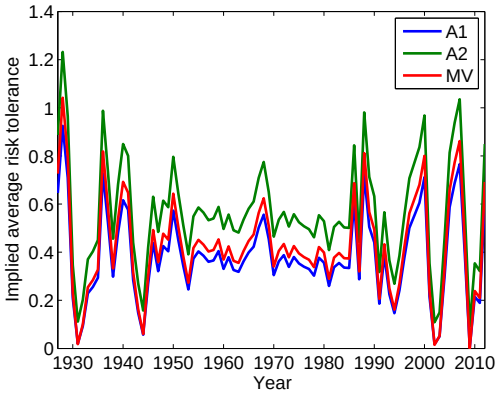

Figure 9: Time series of the implied average risk tolerance $\bar{\tau}_t = \text{top1}_t \tau_H + (1 - \text{top1}_t)\tau_L$.

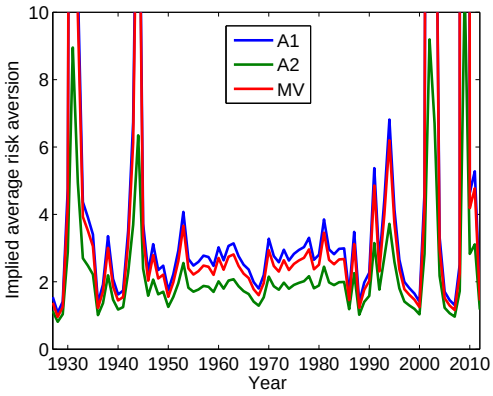

Figure 10: Time series of the implied average risk aversion $1/\bar{\tau}$.

4.3 Two-pass regression

Every asset pricing model is necessarily misspecified, and therefore the standard inference under the null hypothesis may be misleading. There is now a growing literature that documents that when the model is misspecified and contains a factor that is uncorrelated with asset returns ('useless factor'), the factor may spuriously appear to be priced (Kan and Zhang, 1999a,b). We take model misspecification seriously and therefore apply the misspecification-robust two-pass regression of Kan et al. (2013), which is basically the Fama and MacBeth (1973) two-pass regression that gives correct standard errors even under possible model misspecification.

Below, we implement the misspecification-robust two-pass regression following the online appendix of Kan et al. (2013). The models that we consider are the classic CAPM, our two factor with the market return and the top 1% income share, and the Fama and French (1993) three factor model. To impose greater challenges to the asset pricing models, we include the 6 Fama-French portfolios sorted by size and book-to-market value as well as the 5 industry portfolios.[17]

In general, a K-factor beta pricing model specifies that expected excess returns are linear in beta, that is,

$$\mathrm{E}[R_j] - R_f = \gamma_0 + \beta_j' \gamma_1, \tag{4.1}$$

where γ_0 is the difference between the zero-beta rate and the risk-free rate, γ_1 is a K-vector of factor risk premia, and β_j is a K-vector of factor loadings for asset j. In the classic CAPM, since the only factor is the market excess returns, we have $\gamma_1 = \gamma_m$ (m for "market"). In our model, the top 1% income share is also a factor, so $\gamma_1 = (\gamma_m, \gamma_{\text{top1}})$. In the Fama-French three factor model, we have $\gamma_1 = (\gamma_m, \gamma_{\text{smb}}, \gamma_{\text{hml}})$. (smb and hml stand for "small minus big" and "high minus low".)

The estimation of (4.1) proceeds in two steps. In the first pass, we estimate β_j by running a time series regression of excess returns on a constant and factors. In the second pass, we estimate $\gamma = (\gamma_0, \gamma_1)$ by running a cross-sectional

[17] We limit our analysis to these 11 portfolios because consistent estimation requires that the sample size (number of time periods) is much larger than the number of portfolios. Since we use annual data, the sample size is $T = 86$.

regression of average excess returns on a constant and betas. See Kan et al. (2013) for details on how to calculate standard errors.

Table 5 shows the intercept and the betas estimated from the first-pass regressions of the excess returns of Fama-French portfolios on the top 1% income share and the market excess return. We draw two main conclusions from the results. First, consistent with the covariance pricing formula (2.11), both the market return and the top 1% share are statistically significant in explaining the returns for most of the portfolios. In particular, the market return is significant in all six regressions (the classic CAPM result), and the top 1% is significant for small stocks.

Second, the signs of the top 1% coefficients are ordered exactly as the covariance pricing formula would have us expect. As is well known, small stocks have higher average returns than do big stocks, high book-market stocks have higher average returns than do low book-market stocks, and, unsurprisingly, `sh` has a much higher average return than does `bl` (19% vs. 11%). As we see in Table 5, the top 1% coefficients range from 0.645 for `bl` to -5.71 for `sh`. The `sl` coefficient is slightly smaller than the `sm` one, but the coefficients are otherwise ordered as expected. Exactly as in our model (assuming the rich are less risk averse), stocks negatively correlated with the top 1% have high average returns, and stocks positively correlated with the top 1% have low average returns.

Table 5: Time series regressions of Fama-French benchmark portfolios on factors.

Regressors	Portfolio					
	bl	bm	bh	sl	sm	sh
Constant	-0.0717	0.571	1.32	-0.626	2.87*	4.32**
	(0.543)	(0.705)	(1.05)	(1.71)	(1.49)	(1.83)
Market	0.956***	0.982***	1.23***	1.40***	1.27***	1.38***
	(0.0317)	(0.0767)	(0.0921)	(0.104)	(0.107)	(0.0968)
Top 1%	0.645	0.0905	-0.819	-3.96**	-3.65***	-5.71***
	(0.518)	(0.612)	(0.997)	(1.81)	(1.07)	(1.57)
R^2	.936	.891	.832	.776	.819	.775

Newey-West standard errors in parentheses ($k = 4$)
Sample: 1927–2012
***, **, and *: significant at 1%, 5%, and 10%

Table 6 shows the result of the second-pass cross-sectional regression for each model, using the 6 Fama-French portfolios and the 5 industry portfolios as test assets. CAPM, FF3, and Top 1% refer to the classic CAPM, the Fama-French three factor model, and our two factor model. The t-ratios are calculated using the Kan-Robotti-Shanken model misspecification-robust standard errors. The classic CAPM shows a poor fit, with an R-squared of 0.437. As is well-known, the Fama-French three factor model gives a much better fit, with an R-squared of 0.748. The risk premia on the SMB and HML factors are significant but the market return is insignificant. Our two factor model with top 1% income share performs comparably well as the Fama-French three factor model, with an R-squared of 0.668. As expected, the risk premium on the top 1% income share is negative and significant. However, the risk premium on the market return has the wrong sign and it is insignificant.

Table 6: Second-pass estimation of risk premia. FF3: Fama-French three factor model; Top 1%: two factor model with market return and top 1% income share.

Model	CAPM	FF3			Top 1%	
Risk premium	$\hat{\gamma}_m$	$\hat{\gamma}_m$	$\hat{\gamma}_{\text{smb}}$	$\hat{\gamma}_{\text{hml}}$	$\hat{\gamma}_m$	$\hat{\gamma}_{\text{top1}}$
Estimate	7.36*	-7.58	3.56*	2.90*	-3.58	-1.04*
t-ratio	1.90	-1.33	1.96	1.95	-0.514	-1.73
R^2	0.437		0.748		0.668	
N	11		11		11	
T	86		86		86	

Sample: 1927–2012
***, **, and *: significant at 1%, 5%, and 10%

Given that few of the macroeconomic factors considered in the literature (such as consumption growth, human capital, etc.) continue to price assets once we take model misspecification into account (see the empirical analysis of Kan et al. (2013)), it is reassuring that our theoretically motivated top income share factor remains significant and gives a comparable performance to the empirically motivated Fama-French factors.

5 Concluding remarks

In this paper we found that the income/wealth distribution is closely connected with stock market returns. When the rich are richer than usual the stock market subsequently performs poorly. To explain this stylized fact, we built a simple general equilibrium model with agents that are heterogeneous in both wealth and attitudes towards risk. We then derived a testable moment condition as well as a new two factor covariance pricing formula. The formula tells us, essentially, that assets positively correlated with the top income share (and thus the average risk tolerance in the economy) command relatively low risk premiums. Our model is a mathematical formulation of Irving Fisher's narrative that booms and busts are caused by changes in the relative wealth of the rich (the "enterpriser-borrower") and the poor (the "creditor, the salaried man, or the laborer"). Overall, we find that our model is broadly consistent with the data.

Could one exploit the predictive power of top income shares to beat the market on average? The answer is probably no since the top income share—which comes from tax return data—is calculated with a substantial lag. One would receive the inequality update too late to act on its asset pricing information. However, our analysis provides a novel positive explanation of both market excess returns over time and the cross section of returns across stocks. We conclude, as decades of macro/finance theory have suggested, that stock market fluctuations are intimately tied to the distribution of wealth, income, and assets.

A Robustness of predictability

Table 7 shows some of the regressions of Table 1, except with five year stock returns. Table 8 explores the robustness of the result that when the top income or wealth share is above trend, subsequent one year excess returns are significantly below average. Column (1) shows the regression for the .1% income share from Piketty and Saez (2003). Compared with the 1%, the relationship is actually much stronger. Column (2) uses the Piketty and Saez (2003) 1% series in which the income rank includes capital gains. This mitigates the relationship but only slightly. Columns (3) and (4) use the *estimated* top wealth share series of Saez and Zucman (2014). The impact of inequality on asset returns is smaller than with the Kopczuk and Saez (2004) wealth series, but the coefficients remain significant.

In the other columns, we detrend the 1% series but without using the HP filter with a smoothing parameter of 100. First, using a smoothing parameter of 10 strengthens the relationship between inequality and asset returns. In column (5), we use the one-sided HP filter with a smoothing of 100. The one-sided HP filter detrends each data point by applying the filter only to the previous data. This method gives slightly weaker results than our baseline regression. In column (7), we estimate and remove two linear trends, a downward one pre-1977 and an upward one post-1977. Doing so weakens the relationship somewhat, but the inequality coefficient remains significant. We detrend using the Kalman filter method outlined in Appendix C. Column (9) uses and AR(1) model for the cyclical component and column (10) is AR(2). (We also tried white noise ($p = 0$) for the cyclical component, but then the trend becomes almost identical to the raw series. Letting $p \geq 3$ is similar to AR(2).) In these regressions, the 1% series remains significant. Since the Kalman filter uses only current and past data, these results show that the look-ahead bias of the HP filter is not severe. Finally, in column (11) we estimate the trend using a ten year moving average. This method, which is also one-sided, yields a slightly weaker but still significant relationship inequality and subsequent excess returns.

Table 7: Regressions of five year excess stock market returns on top income shares and other predictors

	Dependent Variable: t to $t+5$ Excess Stock Market Return						
Regressors (t)	(1)	(2)	(3)	(4)	(5)	(6)	(9)
Constant	38.44	38.52	43.07	152.42	159.70	38.35	177.39
	(8.54)	(8.81)	(10.14)	(45.59)	(52.38)	(8.38)	(60.79)
Top 1%	-25.12***			-20.61**	-19.11**	-10.45♣	-7.67
	(8.15)			(8.07)	(8.18)	(7.76)	(6.91)
Top 1% (no cg)		-28.62**					
		(11.41)					
Top 1% (wealth)			-14.23***				
			(5.20)				
log(P/D10)				-34.50**			-39.48**
				(13.05)			(16.35)
log(P/E10)					-44.92**		
					(18.55)		
CAY						12.35***	11.09***
						(1.96)	(2.79)
Sample	1913-	1913-	1916-	1913-	1913-	1945-	1945-
	-2008	-2008	-2000	-2008	-2008	-2008	-2008
R^2	.20	.16	.08	.30	.31	.38	.50

Newey-West standard errors in parentheses ($k = 4$)
***, **, and * indicate significance at 1%, 5%, and 10% levels (suppressed for constants)
♣ significant at 10% level with OLS standard error

Table 8: Regressions of one year excess stock market returns on top income and wealth shares

	Dependent Variable: t to $t+1$ Excess Stock Market Return									
Regressors (t)	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(9)	(10)	(11)
Constant	6.67	6.67	6.71	6.71	6.68	6.62	6.73	7.15	7.07	7.45
	(1.70)	(1.77)	(1.69)	(1.69)	(1.95)	(1.74)	(1.72)	(1.81)	(1.79)	(1.90)
Top .1%	-15.32***									
	(4.09)									
Top 1% (rank with cg)		-4.85***								
		(1.29)								
Top 1% (SZ wealth)			-2.87**							
			(1.14)							
Top .1% (SZ wealth)				-3.84***						
				(1.29)						
Top 1% (one-sided HP)					-5.48**					
					(2.41)					
Top 1% (HP, $\lambda = 10$)						-8.13***				
						(1.97)				
Top 1% (linear detrending)							-3.82***			
							(1.18)			
Top 1% (Kalman, $p = 1$)								-3.76**		
								(1.83)		
Top 1% (Kalman, $p = 2$)									-3.46**	
									(1.50)	
Top 1% (10 year MA)										-2.54**
										(1.13)
Sample	1913-	1913-	1913-	1913-	1936-	1913-	1913-	1913-	1913-	1922-
	-2012	-2012	-2012	-2012	-2012	-2012	-2012	-2012	-2012	-2012
R^2	.09	.12	.05	.05	.06	.11	.09	.03	.05	.05

Newey-West standard errors in parentheses ($k = 4$)
***, **, and * indicate significance at 1%, 5%, and 10% levels (suppressed for constants)

B Numerical algorithm

This appendix explains how to compute the equilibrium of the numerical examples in Section 2.3 in the general case. Suppose that there are I agents and J risky assets. Interpret the risky assets as constant-returns-to-scale, stochastic savings technologies; let $\mathbf{R} = (R_1, \dots, R_J)$ be the vector of gross returns with expected return $\mu = \mathrm{E}[\mathbf{R}]$ and variance-covariance matrix $\Sigma = \mathrm{E}[(\mathbf{R}-\mu)(\mathbf{R}-\mu)']$. The equilibrium objects are the portfolios $\{\theta_i^*\}_{i=1}^I$ and the risk-free rate R_f.

First we consider the approximation (2.5). By the budget constraint, we get

$$\frac{w_{i1}}{w_{i0}} = \sum_{j=1}^J R_j\theta_j + R_f\left(1 - \sum_{j=1}^J \theta_j\right) = R_f + \langle \mathbf{R} - R_f 1, \theta\rangle .$$

By the approximation of the first-order condition (2.3) with $a_i = w_{i0}$, we get

$$\begin{aligned} &\mathrm{E}[(\tau_i + 1 - (R_f + \langle \mathbf{R} - R_f 1, \theta\rangle))(\mathbf{R} - R_f 1)] = 0 \\ \iff &(\tau_i + 1 - R_f)(\mu - R_f 1) - \mathrm{E}[(\mathbf{R} - R_f 1)(\mathbf{R} - R_f 1)']\theta = 0 \\ \iff &\theta_i^* = (\tau_i + 1 - R_f)[\Sigma + (\mu - R_f 1)(\mu - R_f 1)']^{-1}(\mu - R_f 1). \end{aligned}$$

This equation shows that the mutual fund theorem holds. Taking the weighted average across agents and using market clearing, it follows that

$$\theta_m = (\bar\tau + 1 - R_f)[\Sigma + (\mu - R_f 1)(\mu - R_f 1)']^{-1}(\mu - R_f 1),$$

where θ_m is the market portfolio and $\bar\tau$ is the average risk tolerance. We can solve for the equilibrium by the shooting algorithm: given a risk-free rate R_f, we compute the market portfolio θ_m, raise the interest rate if $\sum_j \theta_{mj} > 1$ and cut otherwise. Then iterate until we get $\left|\sum_j \theta_{mj} - 1\right| < \varepsilon$, where ε is the error tolerance. We can use the risk-free rate computed from the mean-variance model as an initial guess: by (2.7) and $1'\theta_m = 1$, we obtain

$$R_f^0 = \frac{1'\Sigma^{-1}\mu - 1/\bar\tau}{1'\Sigma^{-1}1}.$$

Next we consider the approximation (2.4). Using the approximation of the first-order condition (2.3) with $a = a_i = \mathrm{E}[w_{i1}]$ and noting that

$$\frac{a_i}{w_{i0}} = \frac{\mathrm{E}[w_{i1}]}{w_{i0}} = R_f + \langle \mu - R_f 1, \theta\rangle ,$$

we obtain

$$\begin{aligned} &\mathrm{E}[(\tau_i(R_f + \langle \mu - R_f 1, \theta\rangle) - \langle \mathbf{R} - \mu, \theta\rangle)(\mathbf{R} - R_f 1)] = 0 \\ \iff &(\mathrm{E}[(\mathbf{R} - R_f 1)(\mathbf{R} - \mu)'] - \tau_i(\mu - R_f 1)(\mu - R_f 1)')\theta = \tau_i R_f(\mu - R_f 1) \\ \iff &\theta_i^* = \tau_i R_f[\Sigma - \tau_i(\mu - R_f 1)(\mu - R_f 1)']^{-1}(\mu - R_f 1). \end{aligned}$$

In this case the mutual fund theorem does not hold, but we can still compute the equilibrium by the shooting algorithm. Starting from some R_f, for each agent we compute the optimal portfolio θ_i^* by the above formula. Then we compute the market portfolio $\theta_m = \sum_i w_{i0}\theta_i^* / \sum_i w_{i0}$ and raise or cut the interest rate according as $\sum_j \theta_{mj} \gtrless 1$.

Solving the exact model is similar, except that for each agent we have to numerically solve the optimal portfolio problem

$$\max_{\theta} \mathrm{E}[u_i((R_f + \langle \mathbf{R} - R_f 1, \theta \rangle) w_{i0})].$$

Letting $V(\theta)$ be the objective function, if the functional form of u_i' and u_i'' are explicitly known, we can solve this problem by the Newton algorithm since the gradient and the Hessian can be computed as

$$\begin{aligned}
\nabla V(\theta) &= \mathrm{E}[u_i'((R_f + \langle \mathbf{R} - R_f 1, \theta \rangle) w_{i0})(\mathbf{R} - R_f 1)], \\
\nabla^2 V(\theta) &= \mathrm{E}[u_i''((R_f + \langle \mathbf{R} - R_f 1, \theta \rangle) w_{i0})(\mathbf{R} - R_f 1)(\mathbf{R} - R_f 1)'],
\end{aligned}$$

respectively.

C Kalman filter

This appendix explains how we detrend the top income/wealth share using the Kalman filter.

Let y_t be the observed top income/wealth share data at time t. Let

$$y_t = g_t + u_t, \tag{C.1}$$

where g_t is the trend and u_t is the cyclical component. We conjecture that the trend is an I(2) process and the cycle is an AR(p) process, so

$$(1 - L)^2 g_t = \epsilon_t, \qquad \epsilon_t \sim i.i.d.\ N(0, \sigma_\epsilon^2), \tag{C.2a}$$

$$\phi(L) u_t = w_t, \qquad w_t \sim i.i.d.\ N(0, \sigma_w^2), \tag{C.2b}$$

where L is the lag operator and

$$\phi(z) = 1 - \phi_1 z - \cdots - \phi_p z^p$$

is the lag polynomial for the autoregressive process. For concreteness, assume $p = 1$ so $\phi(z) = 1 - \phi_1 z$. Then (C.1) and (C.2) can be written as

$$\begin{bmatrix} g_t \\ g_{t-1} \end{bmatrix} = \begin{bmatrix} 2 & -1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} g_{t-1} \\ g_{t-2} \end{bmatrix} + \begin{bmatrix} \epsilon_t \\ 0 \end{bmatrix}, \tag{C.3a}$$

$$y_t = \phi_1 y_{t-1} + g_t - \phi_1 g_{t-1} + w_t. \tag{C.3b}$$

Letting $\xi_t = (g_t, g_{t-1})'$, $v_t = (\epsilon_t, 0)'$, $x_t = y_{t-1}$, $A = \phi_1$, $F = \begin{bmatrix} 2 & -1 \\ 1 & 0 \end{bmatrix}$, and $H = \begin{bmatrix} 1 & -\phi_1 \end{bmatrix}$, (C.3) reduces to

$$\xi_t = F \xi_{t-1} + v_t, \tag{C.4a}$$

$$y_t = A x_t + H \xi_t + w_t. \tag{C.4b}$$

(C.4a) is the state equation and (C.4b) is the observation equation of the state space model. We can then estimate the model parameters $\phi_1, \sigma_\epsilon^2, \sigma_w^2$ as well as the trend $\{g_t\}$ by maximum likelihood: see Chapter 13 of Hamilton (1994) for details. The extension to general AR(p) model is straightforward.

References

Pierluigi Balduzzi and Tong Yao. Testing heterogeneous-agent models: An alternative aggregation approach. *Journal of Monetary Economics*, 54(2): 369–412, March 2007. doi:10.1016/j.jmoneco.2005.08.021.

Ravi Bansal, Dana Kiku, Ivan Shaliastovich, and Amir Yaron. Volatility, the macroeconomy, and asset prices. *Journal of Finance*, 69(6):2471–2511, December 2014. doi:10.1111/jofi.12110.

Harjoat S. Bhamra and Raman Uppal. Asset prices with heterogeneity in preferences and beliefs. *Review of Financial Studies*, 27(2):519–580, 2014. doi:10.1093/rfs/hht051.

Lawrence Blume and David Easley. If you're so smart, why aren't you rich? belief selection in complete and incomplete markets. *Econometrica*, 74(4): 929–966, July 2006. doi:10.1111/j.1468-0262.2006.00691.x.

Alon Brav, George M. Constantinides, and Christopher C. Geczy. Asset pricing with heterogeneous consumers and limited participation: Empirical evidence. *Journal of Political Economy*, 110(4):793–824, August 2002. doi:10.1086/340776.

Douglas T. Breeden. An intertemporal asset pricing model with stochastic consumption and investment opportunities. *Journal of Financial Economics*, 7(3):265–296, September 1979. doi:10.1016/0304-405X(79)90016-3.

Alessandro Bucciol and Raffaele Miniaci. Household portfolio and implicit risk preference. *Review of Economics and Statistics*, 93(4):1235–1250, 2011. doi:10.1162/REST_a_00138.

Laurent E. Calvet and Paolo Sodini. Twin picks: Disentangling the determinants of risk-taking in household portfolios. *Journal of Finance*, 69(2): 867–906, April 2014. doi:10.1111/jofi.12125.

John Y. Campbell. A variance decomposition for stock returns. *Economic Journal*, 101(405):157–179, March 1991. doi:10.2307/2233809.

John Y. Campbell. Household finance. *Journal of Finance*, 61(4):1553–1604, August 2006. doi:10.1111/j.1540-6261.2006.00883.x.

Christopher D. Carroll. Portfolios of the rich. In Luigi Guiso, Michael Haliassos, and Tullio Jappelli, editors, *Household Portfolios*, chapter 10, pages 389–430. MIT Press, Cambridge, MA, 2002.

Yeung Lewis Chan and Leonid Kogan. Catching up with the Joneses: Heterogeneous preferences and the dynamics of asset prices. *Journal of Political Economy*, 110(6):1255–1285, December 2002. doi:10.1086/342806.

John H. Cochrane. *Asset Pricing*. Princeton University Press, Princeton, NJ, second edition, 2005.

Daniele Coen-Pirani. Effects of differences in risk aversion on the distribution of wealth. *Macroeconomic Dynamics*, 8(5):617–632, 2004. doi:10.1017.S1365100504040052.

Timothy Cogley. Idiosyncratic risk and the equity premium: Evidence from the consumer expenditure survey. *Journal of Monetary Economics*, 49(2): 309–334, March 2002. doi:10.1016/S0304-3932(01)00106-4.

George M. Constantinides. Intertemporal asset pricing with heterogeneous agents without demand aggregation. *Journal of Business*, 55(2):253–267, April 1982.

George M. Constantinides and Darrell Duffie. Asset pricing with heterogeneous consumers. *Journal of Political Economy*, 104(2):219–240, April 1996.

Bernard Dumas. Two-person dynamic equilibrium in the capital market. *Review of Financial Studies*, 2(2):157–188, 1989. doi:10.1093/rfs/2.2.157.

Eugene F. Fama and Kenneth R. French. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3–56, February 1993. doi:10.1016/0304-405X(93)90023-5.

Eugene F. Fama and James D. MacBeth. Risk, return, and equilibrium: Empirical test. *Journal of Political Economy*, 81(3):607–636, May-June 1973.

Irving Fisher. *Introduction to Economic Science*. Macmillan, NY, 1910.

Nicolae Gârleanu and Stavros Panageas. Young, old, conservative, and bold: The implications of heterogeneity and finite lives for asset pricing. *Journal of Political Economy*, 2012. Forthcoming.

John Geanakoplos and Martin Shubik. The capital asset pricing model as a general equilibrium with incomplete markets. *Geneva Papers on Risk and Insurance Theory*, 15(1):55–71, 1990. doi:10.1007/BF01498460.

Christian Gollier. Wealth inequality and asset pricing. *Review of Economic Studies*, 68(1):181–203, January 2001. doi:10.1111/1467-937X.00165.

William M. Gorman. Community preference fields. *Econometrica*, 21(1):63–80, January 1953. doi:10.2307/1906943.

Fatih Guvenen. A parsimonious macroeconomic model for asset pricing. *Econometrica*, 77(6):1711–1750, November 2009. doi:10.3982/ECTA6658.

James D. Hamilton. *Time Series Analysis*. Princeton University Press, Princeton, NJ, 1994.

Lars Peter Hansen and Kenneth J. Singleton. Stochastic consumption, risk aversion, and the temporal behavior of asset returns. *Journal of Political Economy*, 91(2):249–265, April 1983.

Chiaki Hara, James Huang, and Christoph Kuzmics. Representative consumer's risk aversion and efficient risk-sharing rules. *Journal of Economic Theory*, 137 (1):652–672, November 2007. doi:10.1016/j.jet.2006.11.002.

Juan Carlos Hatchondo. A quantitative study of the role of wealth inequality on asset prices. *Federal Reserve Bank of Richmond Economic Quarterly*, 94 (1):73–96, 2008.

John Heaton and Deborah J. Lucas. Evaluating the effects of incomplete markets on risk sharing and asset pricing. *Journal of Political Economy*, 104(3):443–487, June 1996.

Timothy C. Johnson. Inequality risk premia. *Journal of Monetary Economics*, 59(6):565–580, October 2012. doi:10.1016/j.jmoneco.2012.06.008.

Raymond Kan and Chu Zhang. Two-pass tests of asset pricing models with useless factors. *Journal of Finance*, 54(1):203–235, February 1999a. doi:10.1111/0022-1082.00102.

Raymond Kan and Chu Zhang. GMM tests of stochastic discount factor models with useless factors. *Journal of Financial Economics*, 54(1):103–127, October 1999b. doi:10.1016/S0304-405X(99)00033-1.

Raymond Kan, Cesare Robotti, and Jay Shanken. Pricing model performance and the two-pass cross-sectional regression methodology. *Journal of Finance*, 68(6):2617–2649, December 2013. doi:10.1111/jofi.12035.

Narayana R. Kocherlakota and Luigi Pistaferri. Asset pricing implications of Pareto optimality with private information. *Journal of Political Economy*, 117(3):555–590, June 2009. doi:10.1086/599761.

Wojciech Kopczuk and Emmanuel Saez. Top wealth shares in the United States, 1916–2000: Evidence from estate tax returns. *National Tax Journal*, 57(2): 445–487, 2004.

Per Krusell and Anthony A. Smith, Jr. Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106(5):867–896, October 1998. doi:10.1086/250034.

Martin Lettau and Sydney C. Ludvigson. Consumption, aggregate wealth, and expected stock returns. *Journal of Finance*, 56(3):815–849, June 2001. doi:10.1111/0022-1082.00347.

Martin Lettau, Sydney C. Ludvigson, and Jessica A. Wachter. The declining equity premium: What role does macroeconomic risk play? *Review of Financial Studies*, 21(4):1653–1687, 2008. doi:10.1093/rfs/hhm020.

John Lintner. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *Review of Economics and Statistics*, 47(1):13–37, February 1965a. doi:10.2307/1924119.

John Lintner. Security prices, risk, and maximal gains from diversification. *Journal of Finance*, 20(4):587–615, December 1965b. doi:10.1111/j.1540-6261.1965.tb02930.x.

Francis A. Longstaff and Jiang Wang. Asset pricing and the credit market. *Review of Financial Studies*, 25(11):3169–3215, 2012. doi:10.1093/rfs/hhs086.

Robert E. Lucas, Jr. Asset prices in an exchange economy. *Econometrica*, 46 (6):1429–1445, November 1978. doi:10.2307/1913837.

Sydney C. Ludvigson. Advances in consumption-based asset pricing: Empirical tests. In George M. Constantinides, Milton Harris, and René M. Stultz, editors, *Handbook of the Economics of Finance*, volume 2, chapter 12, pages 799–906. Elsevier, Amsterdam, 2013. doi:10.1016/B978-0-44-459406-8.00012-3.

N. Gregory Mankiw. The equity premium and the concentration of aggregate shocks. *Journal of Financial Economics*, 17(1):211–219, September 1986. doi:10.1016/0304-405X(86)90012-7.

Harry Markowitz. Portfolio selection. *Journal of Finance*, 7(1):77–91, March 1952. doi:10.1111/j.1540-6261.1952.tb01525.x.

Rajnish Mehra and Edward C. Prescott. The equity premium: A puzzle. *Journal of Monetary Economics*, 15(2):145–161, March 1985. doi:10.1016/0304-3932(85)90061-3.

Thomas Piketty and Emmanuel Saez. Income inequality in the United States, 1913–1998. *Quarterly Journal of Economics*, 118(1):1–41, 2003. doi:10.1162/00335530360535135.

Emmanuel Saez and Gabriel Zucman. Wealth inequality in the United States since 1913: Evidence from capitalized income tax data. NBER Working Paper, 2014. URL `http://www.nber.org/papers/w20625`.

Alvaro Sandroni. Do markets favor agents able to make accurate prediction? *Econometrica*, 68(6):1303–1341, November 2000. doi:10.1111/1468-0262.00163.

William F. Sharpe. Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance*, 19(3):425–442, September 1964. doi:10.1111/j.1540-6261.1964.tb02865.x.

Robert J. Shiller. Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review*, 71(3):421–436, June 1981.

James Tobin. Liquidity preference as behavior towards risk. *Review of Economics and Statistics*, 25(2):65–86, February 1958.

Alexis Akira Toda and Kieran Walsh. The double power law in consumption and implications for testing Euler equations. *Journal of Political Economy*, 2014. URL `http://ssrn.com/abstract=2319454`. Forthcoming.

Annette Vissing-Jørgensen. Towards and explanation of household portfolio choice heterogeneity: Nonfinancial income and participation cost structures. NBER Working Paper, 2002. URL `http://www.nber.org/papers/w8884`.

Jiang Wang. The term structure of interest rates in a pure exchange economy with heterogeneous investors. *Journal of Financial Economics*, 41(1):75–110, May 1996. doi:10.1016/0304-405X(95)00854-8.

Ivo Welch and Amit Goyal. A compehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies*, 21(4):1455–1508, 2008. doi:10.1093/rfs/hhm014.