Journal of Monetary Economics 59 (2012) 319–335

Contents lists available at SciVerse ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jme

Aggregate implications of micro asset market segmentation

Chris Edmond [a,*], Pierre-Olivier Weill [b]

[a] Department of Economics, University of Melbourne, Parkville, VIC 3010, Australia
[b] UCLA, CEPR and NBER, United States

ARTICLE INFO

Article history:
Received 28 January 2011
Received in revised form
9 March 2012
Accepted 23 March 2012
Available online 9 April 2012

ABSTRACT

An extensive empirical literature finds that micro asset markets are segmented from one another. We develop a consumption-based asset pricing model to quantify the aggregate implications of a financial system comprised of many such segmented micro asset markets. We specify exogenously the level of segmentation that determines how much idiosyncratic risk traders bear in their micro market and calibrate the segmentation to match facts about systematic and idiosyncratic return volatility. In our benchmark model traders bear 30% of their idiosyncratic risk, the unconditional aggregate equity premium is 2.4% annual, and the welfare costs of segmentation are substantial, 1.8% of lifetime consumption.

1. Introduction

Do market-specific frictions matter for aggregate asset prices? An extensive empirical literature finds that micro asset markets are *segmented* from one another, in the sense that “local factors”, specific to the market under consideration, help to explain asset prices in that particular market (Collin-Dufresne et al., 2001; Gabaix et al., 2007, for example). These empirical segmentation patterns are commonly interpreted as evidence that contractual constraints, between financial firms, their employees, and their outside investors, create what Shleifer and Vishny (1997) called *limits to arbitrage*. But these analyses give no clear sense of whether such segmentation matters in the aggregate. To address this question, we construct a consumption-based asset pricing model from a collection of segmented micro asset markets. Our approach is deliberately macro: the model does not address particular features of any specific asset market, but can spell out precisely the aggregate implications of the market segmentation frictions.

In our benchmark model, there are many durable risky assets. Each type of asset is traded in its own specialized market. If these risky assets could be frictionlessly traded across markets, all idiosyncratic market-specific risk would be diversified away and traders would be exposed only to aggregate risk. This full risk sharing is prevented by imposing, exogenously, the following pattern of market-specific segmentation frictions: for each market m, an exogenous fraction λ_m of the expense of purchasing assets in that market must be borne by traders specialized in that market. In return, these traders receive λ_m of the benefit, i.e., of the dividends and resale price of assets sold in that market. In equilibrium, the parameter λ_m determines the fraction of *non-tradeable* idiosyncratic risk in market m. When $\lambda_m = 0$ all idiosyncratic risk can be traded and traders are fully diversified. When $\lambda_m = 1$ traders cannot trade away their idiosyncratic risk and instead simply consume the dividends from the asset in their specific market.

Our setup is made tractable by following Lucas (1990) in assuming that investors can pool the *tradeable* idiosyncratic risk within a large family. In equilibrium, the “state price” of a unit of consumption in each market m is a weighted average

* Corresponding author. Tel.: +61 3 8344 9733; fax: +61 3 8344 6899.
E-mail addresses: cedmond@unimelb.edu.au (C. Edmond), poweill@econ.ucla.edu (P.-O. Weill).

0304-3932/$ - see front matter
http://dx.doi.org/10.1016/j.jmoneco.2012.03.006

of the marginal utility of consumption in that market (with weight λ_m) and a term that reflects the cross-sectional average marginal utility of consumption (with weight $1-\lambda_m$). In the special case where $\lambda_m = 0$ for all markets m, then the state price of consumption is equal across markets and equal to the marginal utility of the aggregate endowment. This economy thus collapses to the standard Lucas (1978) model where asset prices depend only on aggregate consumption risk. By contrast, when $\lambda_m > 0$, asset prices also depend on the amount of idiosyncratic consumption risk ultimately borne by traders. This idiosyncratic consumption risk is determined jointly by (i) the level and cross-sectional variation of segmentation frictions λ_m, and (ii) the distribution of idiosyncratic volatilities across markets.

We start by calibrating a special case of the general model where $\lambda_m = \lambda$ for all markets. The parameters governing the aggregate endowment process and preferences are standard: independently and identically distributed (IID) lognormal aggregate endowment growth, time- and state-separable expected utility preferences with constant relative risk aversion $\gamma = 4$. The parameters governing the distribution of individual endowments and the single λ are used to simultaneously match the *systematic* return volatility of a well-diversified market portfolio and key time-series properties of an individual stock's *total* return volatility (see Goyal and Santa-Clara, 2003; Bali et al., 2005). This procedure yields segmentation of approximately $\lambda = 0.30$. This model generates a sizeable unconditional equity premium, some 2.4% annual. However, as is familiar from many asset pricing models with expected utility preferences and trend growth, the model has a risk-free rate that is too high and too volatile.

This benchmark model is then extended by allowing for multiple types of market segmentation λ_m, which generates cross-sectional differences in stock return volatilities. This motivates us to pick values for λ_m in order to match the volatilities of portfolios sorted on measures of idiosyncratic volatility, as documented by Ang et al. (2006).

Our main finding is that aggregation matters: with cross-sectional variation in λ_m, the model needs an *average* amount of segmentation of approximately $\overline{\lambda} = 0.10$ to hit our targets, only one-third that of the single λ model. Moreover, this version of the model delivers essentially the same aggregate asset pricing implications as the single λ benchmark despite having only about one-third the average amount of segmentation. The characteristics of the micro markets in this disaggregate economy are quite distinct: some 50% of the aggregate market by value has a λ_m of approximately zero, with the amount of segmentation rising to a maximum of $\lambda_m = 0.37$ for about 2% of the aggregate market by value. We also find that *dispersion* in the amount of segmentation has significant implications. In particular, while the average amount of segmentation is lower in the multiple λ_m model, the *welfare costs* of segmentation are actually larger than in the single λ model. The welfare cost of segmentation is a convex function of λ_m so that, other things equal, an increase in the dispersion of segmentation increases the welfare cost. For the single λ model (with no dispersion), the welfare costs of segmentation are about 1.8% of lifetime consumption. By contrast, for the multiple λ_m model the welfare costs rise to about 3% of lifetime consumption, even though the average amount of segmentation is only one-third that of the single λ model.

To assist in interpreting our results, we compare our segmented markets model to an otherwise similar *incomplete markets* model. Like a standard incomplete markets model, our segmented markets model features uninsured idiosyncratic risk. This risk is priced in the segmented markets model, but it is not priced in the incomplete markets counterpart. As a consequence, idiosyncratic risk leads to a significant aggregate risk premium in the segmented markets model but has no such implications in the incomplete markets model.

1.1. Market frictions in the asset pricing literature

Traditionally, macroeconomists have taken the view that frictions in financial intermediation or other asset trades are small enough to be neglected. In particular, early contributions, such as Lucas (1978) and Breeden (1979), characterize equilibrium asset prices using frictionless models. The quantitative limitations of plausibly calibrated traditional asset pricing models were highlighted by the "equity premium" and "risk-free rate" puzzles of Mehra and Prescott (1985) and Weil (1989).

Since then an extensive literature has attempted to explicitly incorporate market frictions in an attempt to rationalize these and related asset pricing puzzles.[1] Models introducing market frictions have tended to follow one of two approaches. One part of the financial economics literature followed deliberately micro-market approaches, focusing on the impact of specific frictions in specific financial markets. This micro-markets approach is transparent and leads to precise implications but does not lead to any clear sense of whether or why micro asset market frictions matter in the aggregate. Moreover, these models are typically not well integrated with the standard consumption-based asset pricing framework. Others have taken an unabashedly aggregate approach, with some financial friction faced by a representative intermediary (see, e.g., Aiyagari and Gertler, 1999; Kyle and Xiong, 2001; Vayanos, 2005; He and Krishnamurthy, 2010a,b) or by households (see, among others, Heaton and Lucas, 1996; Chien et al., 2011; Pavlova and Rigobon, 2008). The friction "stands in" for a diverse array of real-world micro frictions facing intermediaries and households. In these macro models, financial intermediaries often bear disproportionate amounts of *aggregate* risk, but this implication is inconsistent with the empirical literature on market segmentation, which emphasizes instead that intermediaries bear disproportionate amounts of "local" or *idiosyncratic* risk.

[1] See for example Aiyagari and Gertler (1991), He and Modest (1995) and Luttmer (1996, 1999) for the quantitative evaluation of asset pricing models with trading frictions.

Our approach takes a middle course. Starting from a model that is consistent with intermediaries bearing too much local risk, we work out the aggregation problem. With the aggregation problem solved, our stylized model of a collection of micro-markets that together form a financial system can then be embedded into an otherwise standard asset-pricing model. In a sense, our model can be viewed as a multiple market version of a *limited participation* model of asset prices where agents are restricted in their ability to participate in asset trade. Important early contributions to this approach include Mankiw and Zeldes (1991), Saito (1995) and Basak and Cuoco (1998). State of the art contributions to this literature include Gomes and Michaelides (2008), Guvenen (2009) and Chien et al. (2011).

Section 2 presents the model and shows how to compute equilibrium asset prices. Section 3 calibrates a special case of the model with a single type of market segmentation and Section 4 shows that this model can generate a sizeable equity premium. Section 5 discusses how the equilibrium in our model can be obtained by traders individually optimizing subject to constraints on asset trade both in their specialized market and in other markets. This section also explains how our model relates to standard incomplete markets models. Section 6 extends our benchmark model by allowing for multiple types of market segmentation and calculates the welfare costs of segmentation.[2]

2. Model

Market structure and endowments. The model is a variant on the pure endowment asset pricing models of Lucas (1978), Breeden (1979) and Mehra and Prescott (1985). Time is discrete and denoted $t \in \{0,1,2,\ldots\}$. There are many distinct micro asset markets indexed by $m \in [0,1]$. Each market m is specialized in trading a single type of durable asset with supply normalized to one. Each period the asset produces a stochastic realization of a non-storable dividend $y_{m,t} > 0$. The aggregate endowment available to the entire economy is $y_t \coloneqq \int_0^1 y_{m,t}\,dm$. The aggregate endowment $y_t > 0$ follows an exogenous stochastic process, described in detail below. Conditional on all realized aggregate variables, the endowments $y_{m,t}$ are independently and identically distributed (IID) across markets.

Preferences. We follow Lucas (1990) and use a representative family construct to provide consumption insurance beyond our market-segmentation frictions. The single representative family, which is initially endowed with the entire supply of assets, consists of many, identical, traders who are specialized in particular asset markets. The period utility for the family is $U(c_t) \coloneqq \int_0^1 u(c_{m,t})\,dm$, where $u : \mathbb{R}^+ \to \mathbb{R}$ is a standard increasing concave utility function. Intertemporal utility for the family has the standard time- and state-separable form, $\mathbb{E}_0[\sum_{t=0}^{\infty} \beta^t U(c_t)]$, with constant time discount factor β. The crucial role of the representative family is to eliminate the wealth distribution across markets as an additional endogenous state variable (see, e.g., Alvarez et al., 2002).

Segmentation frictions. We interpret the representative family as a partially integrated financial system. Each trader in market m works at a specialized trading desk that deals in the asset specific to that market (Fig. 1 illustrates). Traders in market m are assumed to bear an exogenous fraction $\lambda_m \in [0,1]$ of the expense of trading in that market and in return receive λ_m of the benefit. The remaining $1-\lambda_m$ of the expense and benefit of trading in that market is shared between family members.

More precisely, given segmentation parameter λ_m, the period budget constraint facing a representative trader in market m is

$$c_{m,t} + \lambda_m p_{m,t} s_{m,t} + (1-\lambda_m) p^F_{t,t} \le \lambda_m (p_{m,t} + y_{m,t}) s_{m,t-1} + (1-\lambda_m)(p^F_{t-1,t} + y^F_t), \tag{1}$$

where $p_{m,t}$ is the ex-dividend price of a share in the asset in market m while $s_{m,t}$ represents share holdings in that asset.

Family accounting. As can be seen from the budget constraint (1), a trader in market m holds directly a number $\lambda_m s_{m,t}$ of shares of asset m. The collection of remaining shares, $(1-\lambda_n)s_{n,t}$ for all $n \in [0,1]$, is collectively held by all family members in a *family portfolio*. The expense and benefit of trading this family portfolio are divided among family members in a manner summarized by the two terms $(1-\lambda_m)p^F_{t,t}$ and $(1-\lambda_m)(p^F_{t-1,t} + y^F_t)$ in the budget constraint. Specifically, the term $(1-\lambda_m)p^F_{t,t}$ on the left-hand side means that the trader in market m is asked to contribute $1-\lambda_m$ of the expense of acquiring the family portfolio this period (ex-dividend). Symmetrically, the term $(1-\lambda_m)(p^F_{t-1,t} + y^F_t)$ on the right-hand side means that the trader receives $1-\lambda_m$ of the benefit from the family portfolio acquired last period (cum-dividend). Thus, a balanced family budget requires that

$$\int_0^1 (1-\lambda_m) p^F_{t,t}\,dm = \int_0^1 (1-\lambda_n) p_{n,t} s_{n,t}\,dn. \tag{2}$$

In words, the total value of all family members' contributions to the family portfolio (the left-hand side) has to equal the total asset value of the family portfolio (the right-hand side). Defining $\overline{\lambda} \coloneqq \int_0^1 \lambda_m\,dm$, we can rewrite this accounting identity as

$$p^F_{t,t} = \int_0^1 \frac{1-\lambda_n}{1-\overline{\lambda}} p_{n,t} s_{n,t}\,dn. \tag{3}$$

[2] Various extensions and further computational details are given in a Supplementary Appendix, Edmond and Weill (2012), available online from the journal's website.

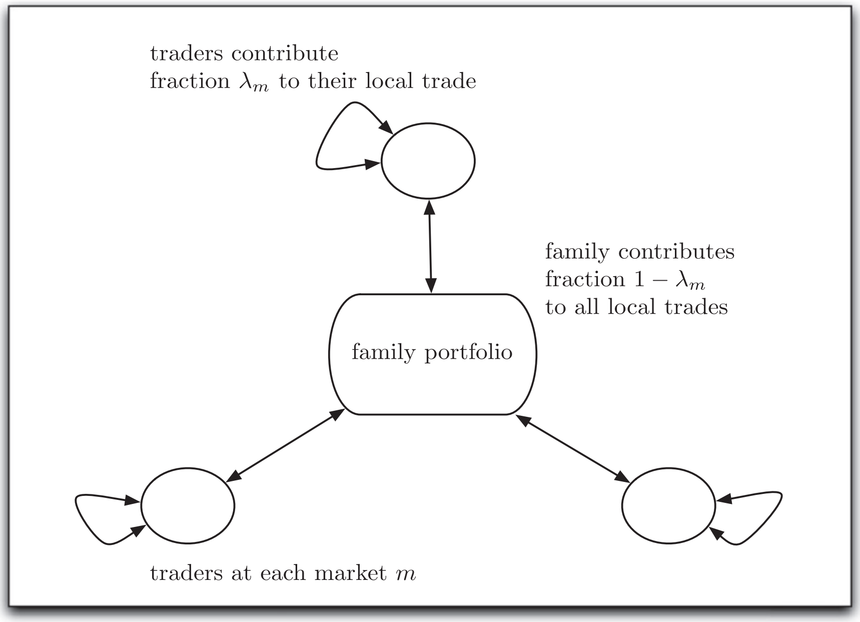

Fig. 1. Segmentation frictions. There are many markets $m \in [0,1]$. Traders at each market bear fraction λ_m of the expense of their trades and share the remaining fraction $1-\lambda_m$ of the expense with all other traders through a *family portfolio*.

Similarly, $\int_0^1 (1-\lambda_m)(p^F_{t-1,t}+y_t)\, dm$ is equal to the cum-dividend value of the remaining shares brought into the period. This yields

$$p^F_{t-1,t}+y^F_t = \int_0^1 \frac{1-\lambda_n}{1-\bar{\lambda}}(p_{n,t}+y_{n,t})s_{n,t-1}\, dn. \tag{4}$$

2.1. Equilibrium

A price path is a sequence $p=\{p_t\}_{t=0}^{\infty}$, adapted to agents' information. Each element of the sequence, $p_t:[0,1]\to\mathbb{R}^+$, is a measurable function mapping each asset $m\in[0,1]$ into its time-t price, $p_{m,t}$. Given a price path, the family maximizes its intertemporal utility by choosing an adapted consumption and asset holding plan, $(c,s)=\{c_t,s_t\}_{t=0}^{\infty}$, where $c_t:[0,1]\to\mathbb{R}^+$ and $s_t:[0,1]\to\mathbb{R}$ are measurable functions specifying $c_{m,t}$ and $s_{m,t}$ in each asset market $m\in[0,1]$. The maximization is subject to the collection of budget constraints (1), one for each $m\in[0,1]$, the accounting identities for the family portfolio, (3) and (4), and takes as given the initial distribution of asset holdings, $s_{m,-1}=1$ for all $m\in[0,1]$.

An *equilibrium* of this economy is a consumption and asset holding plan, (c, s), and a price path, p, such that (i) (c, s) solves the family's problem given p, and (ii) asset markets clear, i.e., $s_{m,t}=1$ for all $m\in[0,1]$ and $t\in\{0,1,2,\ldots\}$.

Equilibrium allocation. Before solving for asset prices, we provide the equilibrium allocation of consumption across markets. Substituting the accounting identities (3) and (4) into the budget constraint (1) and imposing the equilibrium condition $s_{m,t}=1$ gives

$$c_{m,t}=\lambda_m y_{m,t}+(1-\lambda_m)\int_0^1 \frac{1-\lambda_n}{1-\bar{\lambda}} y_{n,t}\, dn. \tag{5}$$

Since the realized idiosyncratic $y_{n,t}$ are independent of λ_n, an application of the law of large numbers then gives

$$c_{m,t}=\lambda_m y_{m,t}+(1-\lambda_m)y_t. \tag{6}$$

Equilibrium consumption in market m is a weighted average of the idiosyncratic and aggregate endowments with weights reflecting the degree of market segmentation. The parameter λ_m represents the extent to which traders are not fully diversified and hence determines the degree of risk sharing in the economy. If $\lambda_m=0$, traders are fully diversified and will have consumption equal to the aggregate endowment, $c_{m,t}=y_t$. But if $\lambda_m=1$, traders are not at all diversified and simply consume the dividends realized in their specific market, $c_{m,t}=y_{m,t}$.

2.2. Asset pricing

Asset prices are obtained using the first-order conditions for the family's problem. Let $\mu_{m,t}\geq 0$ denote the Lagrange multiplier on (1), the constraint for market m at time t. As shown in Appendix A, the family's Lagrangian can be written as

$$\mathcal{L}=\mathbb{E}_0\left[\sum_{t=0}^{\infty}\beta^t\int_0^1 (u(c_{m,t})+q_{m,t}(p_{m,t}+y_{m,t})s_{m,t-1}-q_{m,t}p_{m,t}s_{m,t}-\mu_{m,t}c_{m,t})\, dm\right], \tag{7}$$

where

$$q_{m,t} \coloneqq \lambda_m \mu_{m,t} + (1-\lambda_m) \int_0^1 \frac{1-\lambda_n}{1-\overline{\lambda}} \mu_{n,t} \, dn \tag{8}$$

is a weighted average of the Lagrange multipliers in market m and the multipliers for other markets with weights reflecting the various degrees of market segmentation. More specifically, $q_{m,t}$ is the marginal value to the family of earning one (real) dollar in market m. The first term in (8) arises because a fraction λ_m goes to the local trader, with marginal utility $\mu_{m,t}$. The second term arises because the remaining fraction is shared among other family members, with marginal utility $\mu_{n,t}$, according to their relative contributions $(1-\lambda_n)/(1-\overline{\lambda})$ to the family portfolio. We refer to $q_{m,t}$ as the *state price* of earning one real dollar in market m.

Just as equilibrium consumption in market m is a weighted average of the idiosyncratic or "local" endowment and aggregate endowment with weights λ_m and $1-\lambda_m$, so too the state price for market m is a weighted average of the idiosyncratic multiplier and an aggregate multiplier with the same weights. To highlight this, define

$$q_t \coloneqq \int_0^1 \frac{1-\lambda_n}{1-\overline{\lambda}} \mu_{n,t} \, dn, \tag{9}$$

so that the market-specific state price can be written as $q_{m,t} = \lambda_m \mu_{m,t} + (1-\lambda_m) q_t$. If any particular market m has $\lambda_m = 0$ then the state price in that market is equal to the aggregate state price $q_{m,t} = q_t$ and is independent of the local endowment realization. If the segmentation parameter is common across markets, $\lambda_m = \lambda$ all m, then q_t is the cross-sectional *average marginal utility* and $q_t = \int_0^1 q_{m,t} \, dm$. More generally, q_t is not a simple average over $\mu_{m,t}$ since different markets have different relative contributions $(1-\lambda_m)/(1-\overline{\lambda})$ to the family portfolio.

The first-order conditions for the family are straightforward. For each $c_{m,t}$ we have $u'(c_{m,t}) = \mu_{m,t}$. Taking derivatives with respect to $s_{m,t}$ then gives the Euler equation

$$p_{m,t} = \mathbb{E}_t \left[\beta \frac{q_{m,t+1}}{q_{m,t}} (p_{m,t+1} + y_{m,t+1}) \right], \tag{10}$$

where the expectation is conditional on the family's information at time t. This is a standard equation, familiar from Lucas (1978), with the crucial distinction being that the stochastic discount factor (SDF), $\beta q_{m,t+1}/q_{m,t}$, is *market-specific*.

Combining the formulas for equilibrium consumption (6), market-specific state prices (8), and the pricing equation (10) provides a mapping from the primitives of the economy (the λ_m, $y_{m,t}$, etc.) into equilibrium asset prices. The standard Lucas (1978) asset prices are obtained in the further special case $\lambda_m = 0$ all m, so that $c_{m,t} = y_t$ all m and $\mu_{m,t} = u'(y_t)$ all m and $q_t = \int_0^1 u'(y_t) \, dn = u'(y_t)$.

2.3. Shadow prices of risk-free bonds

To simplify the presentation of the model, we have not explicitly introduced risk-free assets. But "shadow" bond prices can be computed under the following convention. Let $\pi_{k,t}$ denote the price at time t of a zero-coupon bond that pays one unit of the consumption good for sure at time $t+k \geq 1$, and that is held in the family portfolio. As shown in Appendix A, these bonds would have price

$$\pi_{k,t} = \mathbb{E}_t \left[\beta \frac{q_{t+1}}{q_t} \pi_{k-1,t+1} \right], \tag{11}$$

with $\pi_{0,t} \coloneqq 1$. Bonds are priced by the aggregate state price q_t. The one-period shadow gross risk-free rate is $R_{f,t} \coloneqq 1/\pi_{1,t} = 1/\mathbb{E}_t[\beta q_{t+1}/q_t]$. Although the SDF for bonds $\beta q_{t+1}/q_t$ does not depend on any particular idiosyncratic endowment realization, it does depend on the *distribution* of idiosyncratic endowments and in general is *not* Lucas–Breeden SDF.

3. Calibration

Let the period utility $u(c)$ be constant relative risk aversion (CRRA) with coefficient $\gamma > 0$ so that $u'(c) = c^{-\gamma}$. Let the log aggregate endowment be a random walk with drift, $\log g_{t+1} \coloneqq \log(y_{t+1}/y_t) = \log \overline{g} + \epsilon_{g,t+1}$, where the innovations $\epsilon_{g,t+1}$ are IID normal with mean zero and variance $\sigma_{\epsilon g}^2$. Log market-specific endowments are the log aggregate endowment plus an idiosyncratic term, $\log y_{m,t} \coloneqq \log y_t + \log \hat{y}_{m,t}$, so that market-specific endowments inherit the trend in the aggregate endowment. The log idiosyncratic endowment, $\log \hat{y}_{m,t}$, is conditionally IID normal in the cross-section with mean $-\sigma_t^2/2$ and variance σ_t^2 where σ_t follows a stochastic process specified below. The mean is chosen so that the average in levels is normalized to one, i.e., $\int_0^1 \hat{y}_{m,t} \, dm = 1$.

Idiosyncratic endowment volatility. The cross-sectional standard deviation of the idiosyncratic endowment, σ_t, is an AR(1) process in logs

$$\log \sigma_{t+1} = (1-\phi) \log \overline{\sigma} + \phi \log \sigma_t + \epsilon_{v,t+1}, \quad \epsilon_{v,t+1} \sim \text{IID} \quad \text{and} \quad N(0, \sigma_{\epsilon v}^2), \ \overline{\sigma} > 0. \tag{12}$$

For short, we refer to σ_t as *idiosyncratic endowment volatility*, but note that σ_t itself is an *aggregate* state variable. At any point in time, the idiosyncratic endowment volatility σ_t is the same in all markets m. In a frictionless model ($\lambda_m = 0$ all m), all idiosyncratic risk would be diversified away so that asset prices would be independent of the aggregate state σ_t. In other words, despite aggregate fluctuations in the level of idiosyncratic endowment volatility σ_t, the level of σ_t would not be a priced factor. With segmentation frictions ($\lambda_m > 0$), by contrast, both the level and dynamics of σ_t will affect asset prices.

3.1. Solving the quantitative model

Using Eq. (6), equilibrium consumption in market m can be written as the product of the aggregate endowment y_t and an idiosyncratic component that depends only on the local idiosyncratic endowment $\hat{y}_{m,t}$ and the amount of segmentation

$$c_{m,t} = [1+\lambda_m(\hat{y}_{m,t}-1)]y_t. \tag{13}$$

Similarly, using this expression for consumption and the fact that utility is CRRA allows us to rewrite the local state price from (8) as $q_{m,t} = \theta_{m,t} y_t^{-\gamma}$ where

$$\theta_{m,t} \coloneqq \lambda_m[1+\lambda_m(\hat{y}_{m,t}-1)]^{-\gamma} + (1-\lambda_m)\int_0^1 \frac{1-\lambda_n}{1-\overline{\lambda}}[1+\lambda_n(\hat{y}_{n,t}-1)]^{-\gamma}\,dn. \tag{14}$$

The SDF for market m is then $\beta q_{m,t+1}/q_{m,t} = \beta g_{t+1}^{-\gamma}\theta_{m,t+1}/\theta_{m,t}$. As in Campbell and Cochrane (1999) and in recent papers by Lustig and Van Nieuwerburgh (2005), Piazzesi et al. (2007), Kocherlakota and Pistaferri (2009) and Chien and Lustig (2010), amongst others, the SDF can be written as the product of the usual Lucas–Breden aggregate SDF $\beta g_{t+1}^{-\gamma}$ with a multiplicative "twisting" factor $\theta_{m,t+1}/\theta_{m,t}$. Unlike these papers, however, the twisting factor in our model is *market-specific*. The twisting factor varies over time both because of fluctuations in the local endowment $\hat{y}_{m,t}$ and also because of aggregate fluctuations in the *cross-sectional distribution* of endowments, as determined by the volatility factor σ_t. Appendix E discusses the properties of the twisting factor in further detail.

To solve the model in stationary variables, let $\hat{p}_{m,t} \coloneqq p_{m,t}/y_t$ denote the price-to-aggregate-dividend ratio for market m. Dividing both sides of Eq. (10) by $y_t > 0$ and using $g_{t+1} \coloneqq y_{t+1}/y_t$ this ratio solves the Euler equation

$$\hat{p}_{m,t} = \mathbb{E}_t\left[\beta g_{t+1}^{1-\gamma}\frac{\theta_{m,t+1}}{\theta_{m,t}}(\hat{p}_{m,t+1}+\hat{y}_{m,t+1})\right], \tag{15}$$

which is the standard CRRA equation except for the twisting factor $\theta_{m,t+1}/\theta_{m,t}$. This is a linear integral equation to be solved for the unknown function mapping the state into the price/dividend ratio. As detailed in Appendix B, we solve this integral equation numerically using the methods of Tauchen and Hussey (1991).

3.2. Calibration strategy

The model is calibrated to monthly postwar data. The aggregate endowment is interpreted as per capita real personal consumption expenditure on nondurables and services with $\overline{g} = (1.02)^{1/12}$ set to match an annual 2% growth rate and $\sigma_{\epsilon g} = 0.01/\sqrt{12}$ set to match an annual 1% standard deviation. The discount factor is set to $\beta = (0.99)^{1/12}$ to reflect an annual pure rate of time preference of 1% and the coefficient of relative risk aversion is set to $\gamma = 4$.

For our benchmark calibration we assume that all markets in the economy share the same segmentation parameter, λ. Given the values for preference parameters β,γ and the aggregate endowment growth process $\overline{g},\sigma_{\epsilon g}$ above, values still need to be assigned to this single λ and the three parameters of the cross-sectional endowment volatility process $\overline{\sigma},\phi,\sigma_{\epsilon v}$.

3.3. Calibrating the idiosyncratic volatility process

The crucial consequence of market segmentation is that local traders are forced to bear some idiosyncratic risk. Thus, to explain the impact of market segmentation on risk premia, it is important that our model generates realistic levels of idiosyncratic risk. This leads us to choose the parameters of the stochastic process for idiosyncratic endowment volatility in order to match key features of the volatility of a typical stock return. To see why there is a natural mapping between the two volatilities, observe that the gross return on a stock can be written as $R_{m,t} = g_t(\hat{y}_{m,t}+\hat{p}_{m,t})/\hat{p}_{m,t-1}$. Thus, the volatility of $\hat{y}_{m,t}$ directly affects stock returns through the dividend term of the numerator. It also indirectly affects stock returns through the asset price, $\hat{p}_{m,t}$.

Our statistics on stock return volatility draw on Goyal and Santa-Clara (2003). Their measure of monthly stock volatility is obtained by adding up the cross-sectional stock return dispersion over each day of the previous month. Fig. 2 shows the monthly time series of their measure of the cross-sectional standard deviation of stock returns, as updated by Bali et al. (2005).

The idiosyncratic endowment volatility process is chosen so that our model replicates three key features of this stock return volatility data, namely: (i) the unconditional average return volatility of 16.4% monthly, (ii) the unconditional standard deviation of return volatility 4.17% monthly, and (iii) the AR(1) coefficient of return volatility 0.84 monthly.

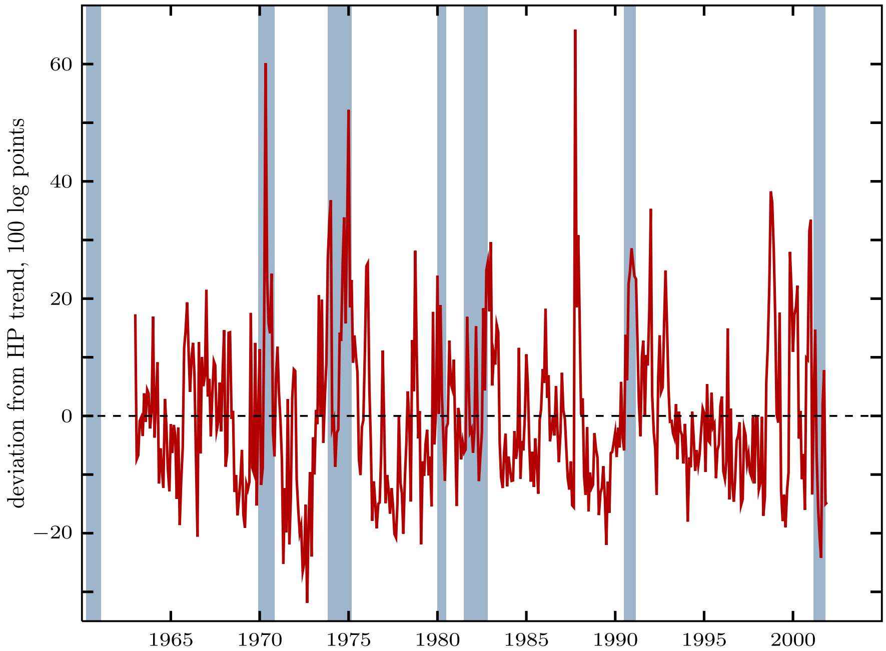

Fig. 2. Cross-sectional standard deviation of stock returns. Cross-sectional standard deviation of CRSP stock returns, monthly (1963:1–2001:12), from Goyal and Santa-Clara (2003) as updated by Bali et al. (2005). Hodrick–Prescott filtered with smoothing parameter 1600×3^4, as recommended by Ravn and Uhlig (2002) for monthly data. NBER recession dates shaded.

These three features are replicated by simultaneously choosing the three parameters governing the stochastic process for endowment volatility: the unconditional average $\overline{\sigma}$, the innovation standard deviation $\sigma_{\epsilon v}$, and the AR(1) coefficient ϕ.

3.4. Calibrating the segmentation parameter

The segmentation parameter λ governs the extent to which local traders can diversify away the return volatility of their local asset. Thus, λ determines the extent to which the volatility factor, σ_t, has an impact on asset prices and creates systematic variation in asset returns. This leads us to identify λ using a measure of systematic volatility, specifically the 4.16% monthly standard deviation of the real value-weighted return of NYSE stocks from CRSP.

To understand how the identification works, recall first what would happen in the absence of market segmentation, $\lambda = 0$. Then, we would be back in Mehra and Prescott model with IID lognormal aggregate endowment growth. As is well known, this model cannot generate realistic amounts of systematic volatility. Specifically, with $\lambda = 0$ the return from a diversified market portfolio is $g_t(1+\overline{p})/\overline{p}$ where $\overline{p} = \beta\mathbb{E}[g^{1-\gamma}]/(1-\beta\mathbb{E}[g^{1-\gamma}])$ is the constant price/dividend ratio for the aggregate market. With our standard parameterization of the preference parameters and aggregate endowment growth, $(1+\overline{p})/\overline{p} \approx 1.0058$ so that the monthly standard deviation of the diversified market portfolio return is approximately the same as the monthly standard deviation of aggregate endowment growth, 0.29% monthly as opposed to 4.16% monthly in the data.

By contrast, with segmentation frictions ($\lambda > 0$), idiosyncratic endowment volatility *creates* systematic volatility. Indeed, because of persistence, high idiosyncratic endowment volatility this month predicts high idiosyncratic endowment volatility next month. Thus in every market m local traders expect to bear more idiosyncratic risk, and, because of risk aversion, the price/dividend ratio $\hat{p}_{m,t} = p_{m,t}/y_t$ has to go down everywhere. Because this effect impacts all stocks at the same time, it endogenously *creates* systematic return volatility. Clearly, the effect is larger if markets are more segmented and traders are forced to bear more idiosyncratic risk. A larger λ will thus result in a larger increase in systematic volatility.

3.5. Calibration results

The calibrated parameters are listed in Table 1. In our benchmark calibration, the level of λ is 0.31. That is, 31% of idiosyncratic endowment risk is non-tradeable. In terms of portfolio weights, $\lambda = 0.31$ also implies that, in a typical market m, a trader invests approximately 31% of his wealth in the local asset and the rest in the family portfolio. Table 2 shows that, with these parameters, the benchmark model matches the target moments exactly.

4. Quantitative examples

Let the gross market return be $R_{M,t+1} := (p_{t+1}+y_{t+1})/p_t$ where $p_t := \int_0^1 p_{m,t}\,dm$ is the ex-dividend value of the market portfolio, y_t is the aggregate endowment, and p_t/y_t is the price/dividend ratio of the market. The shadow gross one period

Table 1
Parameter choices.

Panel A: *Preferences and aggregate endowment growth*

Parameter	Monthly value	Notes
β	0.999	Or 1.004 when calibrated, as below
γ	4	Coefficient relative risk aversion
$\overline{g}$	1.002	Average aggregate growth 2% annual
$\sigma_{\epsilon g}$	0.003	Std dev aggregate growth 1% annual

Panel B: *Segmentation and idiosyncratic endowment volatility*

	Model						
Parameter	Benchmark	Constant	Feedback	$\beta>1$	Data moment		
λ	0.310	0.310	0.310	0.312	Std dev diversified market portfolio return	4.16%	Monthly
$\overline{\sigma}$	0.318	0.318	0.318	0.316	Average cross-section std dev returns	16.40%	Monthly
$\sigma_{\epsilon v}$	0.207	0	0.207	0.205	Time-series std dev cross-section std dev returns	4.17%	Monthly
ϕ	0.784	0	0.785	0.790	AR(1) cross-section std dev returns	0.84	Monthly
η	n/a	n/a	2.513	n/a	Cross-section std dev returns on lagged growth	−0.56	Monthly
β	0.999	0.999	0.999	1.004	Average risk-free rate	1.81%	Annual

The top panel shows our parameters for preferences and aggregate endowment growth. The bottom panel shows our parameters for segmentation and the idiosyncratic endowment volatility process σ_t and the moments in Goyal and Santa-Clara (2003) cross-sectional standard deviation of stock returns data that they are chosen to match. The Benchmark model has a single common segmentation parameter λ and time-varying idiosyncratic endowment volatility σ_t. The Constant σ model sets $\sigma_t=\overline{\sigma}$, i.e., to the Benchmark unconditional mean, for all t. The Feedback model has *counter-cyclical* endowment volatility, with feedback from aggregate growth g_t to volatility σ_t governed by the elasticity η. The $\beta>1$ model chooses β to match the average risk free rate. The Feedback and $\beta>1$ models are re-calibrated, each using an additional moment (as shown) in addition to those moments used for the Benchmark model. For all other cases, β has its benchmark value $\beta=0.999$. See the main text for further details.

Table 2
Fit of calibrated models.

		Model			
Moment	Data	Benchmark	Constant	Feedback	$\beta>1$
Std dev diversified market portfolio return	4.16	4.16	1.01	4.16	4.16
Average cross-section std dev returns	16.40	16.40	16.03	16.35	16.40
Time-series std dev cross-section std dev returns	4.17	4.17	0	4.17	4.17
AR(1) cross-section std dev returns	0.84	0.84	n/a	0.84	0.84
Regression cross-section std dev returns on lagged growth	−0.56	n/a	n/a	−0.56	n/a
Average risk-free rate (annual)	1.81	8.19	9.25	8.19	1.81

Our target moments in the US monthly postwar Goyal and Santa-Clara (2003) cross-sectional standard deviation of stock returns data and their model counterparts. The Benchmark model has a single common segmentation parameter λ and time-varying idiosyncratic endowment volatility σ_t. The Constant σ model sets $\sigma_t=\overline{\sigma}$, i.e., to the Benchmark unconditional mean, for all t. The Feedback model has *counter-cyclical* endowment volatility, with feedback from aggregate growth g_t to volatility σ_t governed by the elasticity η. The $\beta>1$ model chooses β to match the average risk free rate. See the main text for further details.

risk free rate is $R_{f,t} := \mathbb{E}_t[\beta q_{t+1}/q_t]^{-1}$ where q_t is the aggregate state price that determines the price of risk free bonds, as in (11). Implicitly, bonds are priced as if they trade in their own frictionless "$\lambda=0$" market, but the pricing of such bonds takes into account $\lambda>0$ in other asset markets. The *unconditional* equity risk premium is calculated as $\mathbb{E}[R_{M,t+1}-R_{f,t}]$, and similarly for other statistics.

Table 3 shows our model's implications for aggregate returns and price/dividend ratios. The table reports annualized monthly statistics from the model and compares these to annualized monthly returns and to annual price/dividend ratios (annual data for price/dividends are used because of the pronounced seasonality in dividends at the monthly frequency).

4.1. Equity premium

The benchmark model produces an annual equity risk premium of 2.4% annual as opposed to about 5.4% annual in our sample. Clearly this is a much larger equity premium than is produced by a standard Lucas/Mehra and Prescott model. For comparison, that model with risk aversion $\gamma=4$ and IID consumption growth with annual standard deviation of 1% produces an annual equity premium of about 0.04%.

Why is there a large equity premium? Relative to standard consumption-based asset pricing models with time-separable expected utility preferences, our model delivers a large equity premium. Is this a direct consequence our strategy of picking λ in order to match systematic return volatility? No. Model risk premia are generated by *covariances*: no matter

Table 3
Aggregate asset pricing implications of single λ model.

	Moment	Data	Model			
			Benchmark	Constant	Feedback	$\beta>1$
Equity premium	$\mathbb{E}[R_M-R_f]$	5.43	2.43	0.22	2.43	2.12
	$\text{Std}[R_M-R_f]$	14.25	13.27	1.01	13.27	13.34
Sharpe ratio	$\mathbb{E}[R_M-R_f]/\text{Std}[R_M-R_f]$	0.38	0.17	0.20	0.17	0.16
Market return	$\mathbb{E}[R_M]$	7.24	10.62	9.47	10.62	3.93
	$\text{Std}[R_M]$	14.44	14.41	1.01	14.41	14.41
Risk free rate	$\mathbb{E}[R_f]$	1.81	8.19	9.25	8.19	1.81
	$\text{Std}[R_f]$	1.20	5.55	0	5.57	5.41
Price/dividend ratio	$\mathbb{E}[p/y]$ (annual)	34.38	14.10	14.13	14.10	118.97
	$\text{Std}[\log(p/y)]$ (annual)	38.63	20.56	0	20.56	20.90
	$\text{Auto}[\log(p/y)]$ (monthly)	0.99	0.76	n/a	0.76	0.77

Aggregate asset pricing moments in postwar US data. All return data is monthly 1959:1–2007:12 and reported in annualized percent. The stock market index is the value weighted NYSE return from CRSP, and the risk-free return is the 90 day T-bill rate. We obtain real returns after deflating by the CPI from the BLS. Data on price/dividend ratios is annual 1959–2007. To annualize monthly returns we multiply by 12 and to annualize monthly standard deviations we multiply by $\sqrt{12}$.

how much return volatility is fed into a model, the equity risk premia will be zero if the model's SDF is not negatively correlated with return variation.

What, then, is the equity premium from the point of view of aggregate consumption? In our model, if we compute the unconditional average equity premium using the model generated market returns and Lucas–Breeden SDF $\beta g_{t+1}^{-\gamma}$ instead of the true model SDF, then the equity premium is on the order of 0.04% (four basis points) annual rather than the 2.4% annual in the benchmark model. Hence, while aggregate consumption growth does not command a big risk premium, the volatility factor does. To see this, consider the premium implied by the SDF $\beta q_{t+1}/q_t$ where q_t is the aggregate state price that determines the price of risk free bonds. In general this is given by Eq. (9) but with a single common λ it reduces to $q_t = \int_0^1 \mu_{m,t}\, dm = \int_0^1 c_{m,t}^{-\gamma}\, dm$, the cross-section *average marginal utility*. In our benchmark, this SDF implies a premium of 2.05% annual. This comes from the convexity of the marginal utility function: a high σ_t makes consumption highly dispersed across markets so that average marginal utilities are high. At the same time, a high σ_t depresses asset prices in every market, so that the return on the market portfolio is low.

4.2. *Risk free rate and yield curve*

The level of the risk-free rate is high, about 8.2% in the model as opposed to 1.8% in the data. As emphasized by Weil (1989), this comes from the relationship between real interest rates and growth in a deterministic setting with expected utility: high risk aversion means low intertemporal elasticity of substitution so that it takes high real interest rates to compensate for high aggregate growth. With risk, there is an offsetting *precautionary savings* effect that could, in principle, pull the risk-free rate back down to more realistic levels. But in our calibration this precautionary savings effect is small: raising λ from zero to $\lambda=0.31$ lowers the risk free rate by about 1% annual.

In the data, the risk-free rate is smooth and the volatility of the equity premium reflects the volatility of equity returns. In the benchmark model, the risk free-rate is too volatile, about 5.6% annual as opposed to 1.2% annual in the data.

With IID lognormal aggregate growth and CRRA utility, the average yield curve in a standard asset pricing model is flat. But our model generates an increasing and concave average yield curve (see Figure III in Supplementary Appendix). This comes from the relationship between the aggregate state price q_t and volatility σ_t. Since σ_t has positive serial correlation but is not a random walk, its first difference is negatively serially correlated. This negative serial correlation is inherited by the one-period bond pricing SDF $\beta q_{t+1}/q_t$, and this implies that the average yield curve is increasing (Backus and Zin, 1994).

4.3. *Price/dividend ratio*

The benchmark model produces an annual price/dividend ratio of about 14 as opposed to an unconditional average of more like 34 in our sample. Given the large, persistent, swings in the price/dividend ratio in the data, what constitutes success on this dimension is not entirely clear. The model generates too little unconditional volatility in the log price/dividend ratio, some 21% annual as opposed to 39% in the data. Also, the temporal composition of price/dividend volatility differs somewhat between the model and data. The unconditional volatility of the price/dividend ratio in the data comes from large, low-frequency movements whereas in the model it comes from high-frequency movements.

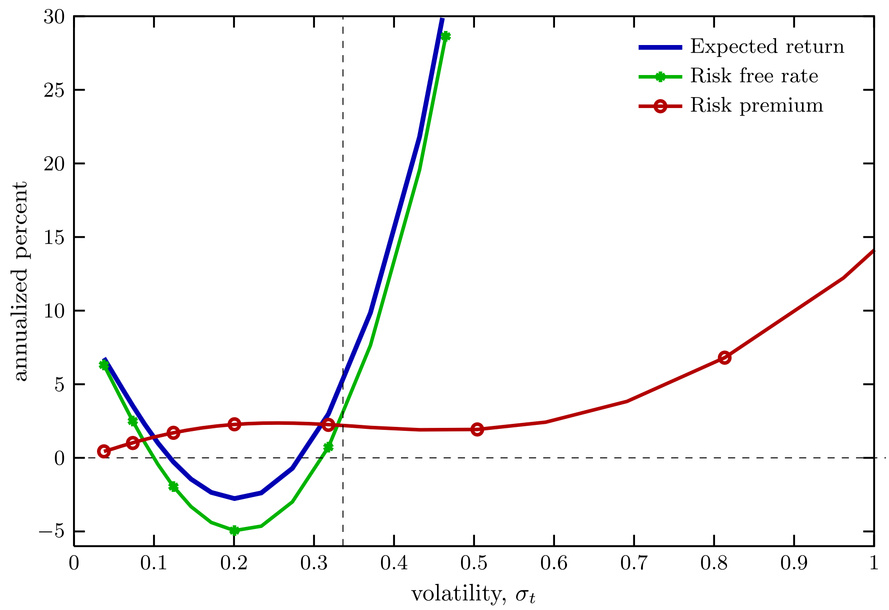

Fig. 3. Conditional returns. The expected market return, risk free rate, and expected excess return (risk premium) as a function of the volatility state σ_t, all expressed in annual terms. The aggregate endowment growth is fixed at is unconditional mean. The vertical dashed line is the unconditional mean $\overline{\sigma}$.

4.4. *Time variation in expected returns*

In a frictionless ($\lambda = 0$) version of our model, all idiosyncratic risk would be diversified and time-variation in the volatility factor σ_t would be irrelevant for asset prices. Since aggregate endowment growth is IID, in that frictionless world, the market price/dividend ratio would be constant as would expected returns and excess returns. Realized returns would inherit the IID property of aggregate endowment growth. In our benchmark model with $\lambda > 0$, however, the volatility factor σ_t is priced. And, since σ_t is persistent, fluctuations in σ_t lead to fluctuations in expected returns and return volatility.

In particular, Fig. 3 shows the expected market return, the risk-free rate and risk premium as a function of σ_t. Except for very low values of σ_t, the expected return and risk-free rate are increasing in σ_t, thus the expected return is relatively high in "bad" aggregate states and relatively low in "good" aggregate states. In this sense, expected returns are countercyclical. However, the risk-free rate is just as cyclical as the market return so that the risk premium is almost a-cyclical. The aggregate market risk premium is significantly countercyclical only for very high values of σ_t, values far above the unconditional mean. Appendix D provides further details.

4.5. *Further discussion*

Constant endowment volatility. Our benchmark model has two departures from a standard consumption-based asset pricing model: (i) segmentation, and (ii) time-varying endowment volatility. To show that both these departures are essential for our results, we solve our model with constant endowment volatility, i.e., $\sigma_t = \overline{\sigma}$ for all t. For this exercise, we fix the volatility at the same level as the unconditional average from the benchmark model $\overline{\sigma} = 0.32$ and keep the level of segmentation at the benchmark $\lambda = 0.31$. Table 2 shows that this version of the model produces essentially the same amount of unconditional cross-sectional stock return volatility as in the data but produces relatively little systematic stock volatility. In particular, systematic stock volatility is only about 1% monthly as opposed to 4% in the data. And recall that, for our preference and aggregate growth parameters, a standard model would imply negligible systematic stock volatility. Thus $\lambda > 0$ is necessary but not sufficient for our model to create systematic stock volatility from idiosyncratic endowment volatility.

Countercyclical endowment volatility. Measures of cross-sectional idiosyncratic risk increase in recessions (Campbell et al., 2001; Storesletten et al., 2004, for example). This cyclicality is also a feature of the cross-sectional standard deviation of returns data from Goyal and Santa-Clara (2003). However, in the benchmark model the stochastic process for the cross-sectional volatility evolves independently of aggregate growth. To see if our results are sensitive to this, we modify the stochastic process in (12) to

$$\log \sigma_{t+1} = (1-\phi)\log \overline{\sigma} + \phi \log \sigma_t - \eta(\log g_t - \log \overline{g}) + \epsilon_{v,t+1}, \tag{16}$$

with $\epsilon_{v,t+1}$ IID normal, as before. If $\eta > 0$, then aggregate growth below trend in period t increases the likelihood that volatility is above trend in period $t+1$. The new parameter η is identified by requiring that, in a monthly regression of the cross-section standard deviation of stock returns on *lagged* aggregate growth, the regression coefficient is -0.56, as it is in the data. The calibrated parameters for this version of the model are shown in Table 1. The elasticity η is 2.5 so aggregate growth 1% below trend tends to increase endowment volatility by 2.5%. The other calibrated parameters are indistinguishable from their benchmark values. The model's implications for asset prices are also very close to the results for the benchmark model. Thus, while the model can be reconciled with the countercyclical behavior of cross-sectional stock volatility, this feature is not necessary for our main results.

Alternative calibration with $\beta > 1$. The level of the risk free rate in our model can be reduced by allowing a pure time discount factor $\beta > 1$. As emphasized by Kocherlakota (1990), since the growth-adjusted discount factor is $\beta g^{1-\gamma}$, for $\gamma > 1$ a value of $\beta > 1$ can still be consistent with finite expected utility. Table 1 presents a version of our model choosing a value of β to match the level of the risk-free rate, with other parameters chosen to match the same moments as before. This gives $\beta = 1.0042$ monthly so that the annual growth-adjusted discount factor is approximately 0.99 and the model risk-free rate is 1.81% annual, on average. The other calibrated values are essentially unchanged and the model's ability to match the target moments is not compromised by the need to also match the risk-free rate. While the model with $\beta > 1$ is able to deliver a lower risk-free rate than our benchmark model, it dramatically increases the average market price/dividend ratio.

5. Individual optimality and trading constraints

In our model, individual traders do not optimize. Instead, the family optimizes on their behalf. In this section we explain two implementation schemes such that the trades dictated by the family are individually optimal. The first implementation uses *portfolio constraints*: an individual trader is constrained to a minimum asset holding if her private valuation of an asset is lower than the market price, and vice versa if her private valuation is higher. The second implementation does the same thing, but using *taxes and subsidies*.

5.1. Implementation with portfolio constraints

Let $V_{m,t}$ denote the equilibrium private valuation of trader m for the asset in that local market and let $V^F_{m,t}$ denote their equilibrium private valuation for the family portfolio. These are given by

$$V_{m,t} := \mathbb{E}_t\left[\beta \frac{u'(c_{m,t+1})}{u'(c_{m,t})}(p_{m,t+1}+y_{m,t+1})\right], \tag{17}$$

$$V^F_{m,t} := \mathbb{E}_t\left[\beta \frac{u'(c_{m,t+1})}{u'(c_{m,t})}(p^F_{t,t+1}+y^F_{t+1})\right], \tag{18}$$

where $c_{m,t}$ denotes the consumption of trader m, $p_{m,t}$ the price of asset m, and $p^F_{t,t}$ the price of the family portfolio in the equilibrium corresponding to the family problem.

With these definitions in mind, we reverse-engineer a simple set of portfolio constraints which make the trades dictated by the family also individually optimal. Consider, for simplicity, the case when $\lambda_m = \lambda$ for all m and suppose that an individual trader can trade his local asset and the family portfolio. Then the trader's sequential budget constraint is

$$c_{m,t}+p_{m,t}s_{m,t}+p^F_{t,t}s^F_{m,t} \le (p_{m,t}+y_{m,t})s_{m,t-1}+(p^F_{t-1,t}+y^F_t)s^F_{m,t-1}. \tag{19}$$

Now assume that the trader faces the following constraints on the quantities of her asset holdings

$$s_{m,t} \ge \lambda \quad \text{if } V_{m,t} \le p_{m,t} \quad \text{and} \quad s_{m,t} \le \lambda \quad \text{if } V_{m,t} \ge p_{m,t},$$

$$s^F_{m,t} \ge 1-\lambda \quad \text{if } V^F_{m,t} \le p^F_{t,t} \quad \text{and} \quad s^F_{m,t} \le 1-\lambda \quad \text{if } V^F_{m,t} \ge p^F_{t,t}. \tag{20}$$

for the local asset and for the family portfolio, respectively. One can immediately verify that the trader's allocation in the family equilibrium solves the problem of an individual trader when faced with these portfolio constraints. The constraints are intuitive. When the private valuation of the trader is below the market price, then the trader wants to lower her holding below λ, and so, to implement the family equilibrium, the trader needs to be confronted with the constraint that $s_{m,t} \ge \lambda$. The opposite is true when the trader's private valuation is above the market price.

Panel A of Fig. 4 illustrates this pattern of binding constraints using our benchmark calibration. Consider for instance the left-side of the graph, when the local endowment realization, $y_{m,t}$, is low. Then consumption $c_{m,t} = \lambda y_{m,t}+(1-\lambda)y_t$ is low as well, implying that the marginal utility of the local trader, $\mu_{m,t}$, is high relative to that of the family, $q_{m,t} = \lambda\mu_{m,t}+(1-\lambda)\int_0^1 \mu_{n,t}\, dn$. As shown in the figure, this means that the local trader has a low private valuation for assets. The portfolio constraint thus prescribes that they should hold a minimum position. The opposite is true when the local endowment realization is high.

5.2. Implementation with taxes and subsidies

The family's trades can also be implemented using taxes and subsidies. The main idea is simply to tax the local trader's asset purchases when their private valuation is high relative to that of the family and to subsidize their purchases when their private valuation is low. Specifically, consider a scheme offering to pay $\tau_{m,t} = 1-V_{m,t}/p_{m,t}$ per real dollar invested in

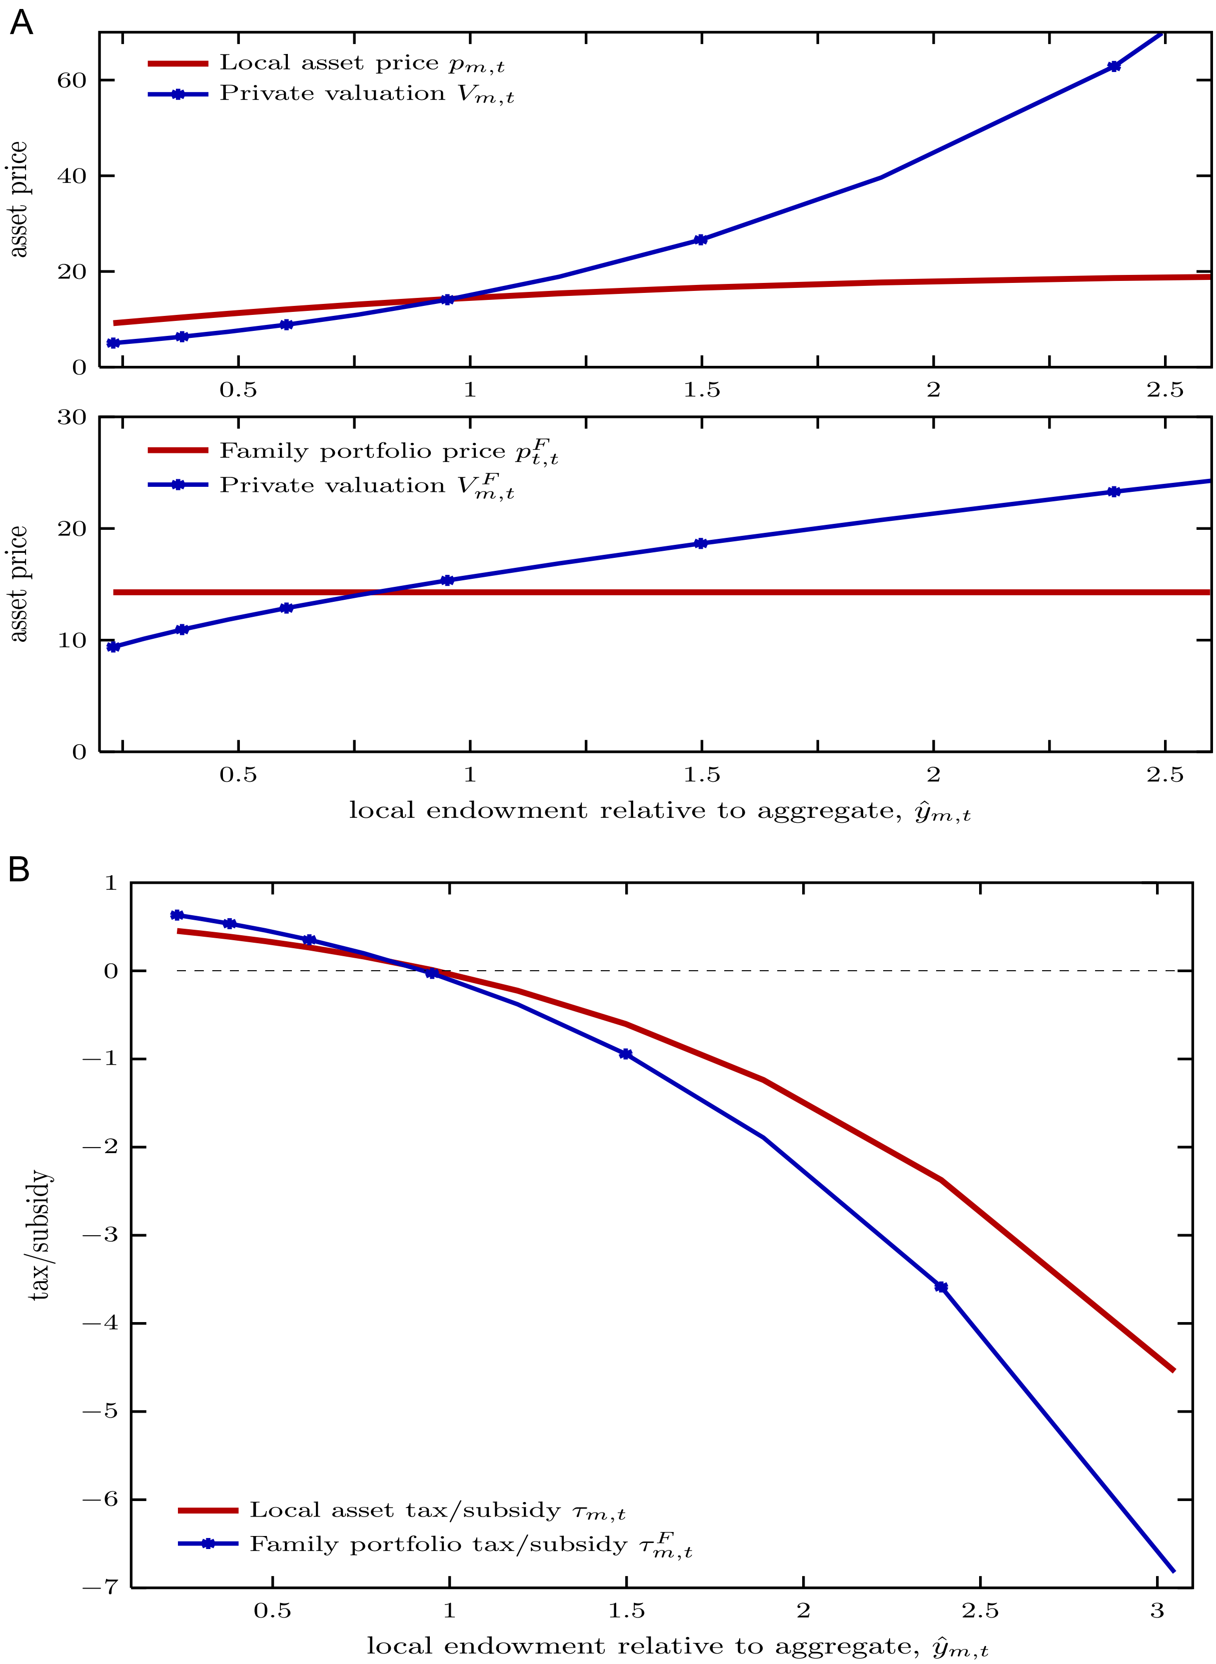

Fig. 4. Implementation of family trades. Panel A: *Implementation with portfolio constraints.* The market price and trader's private valuation for the local asset and for the family portfolio as a function of the local endowment $\hat{y}_{m,t} = y_{m,t}/y_t$, all expressed in annual terms. The aggregate endowment growth and volatility are kept fixed at their unconditional means. Panel B: *Implementation with taxes/subsidies.* The tax/subsidy for the local asset $\tau_{m,t}$ and for the family portfolio $\tau^F_{m,t}$, both as a function of the local endowment $\hat{y}_{m,t} = y_{m,t}/y_t$. The aggregate endowment growth and volatility are kept fixed at their unconditional means.

the local asset and $\tau^F_{m,t} = 1 - V^F_{m,t}/p^F_{t,t}$ per real dollar invested in the family portfolio (if positive, $\tau_{m,t}$ is a subsidy, if negative it is a tax). Panel B of Fig. 4 illustrates the subsidies and taxes using our benchmark calibration.

5.3. Importance of constraining trade in all assets

To implement the family equilibrium, typically there need to be constraints not only on trades in local assets but also on trades in other markets. Moreover, these constraints may prescribe either minimum or maximum holdings.

To highlight the importance of imposing constraints on both kinds of assets, Appendix C studies a version of our model with *only* constraints on local asset trades. Specifically, we consider an incomplete markets version of our model in which aggregate consumption growth and idiosyncratic dividends are independent and IID. Each trader $m \in [0,1]$ is constrained to hold *at least* λ shares of the asset traded in their local market, but faces no constraints on their holdings of assets traded in other markets. Under this portfolio constraint, the model becomes essentially equivalent to the incomplete markets model of Krueger and Lustig (2010), whose predictions are markedly different from those of the segmented markets model. In particular, all local assets are sold at *the same* ex-dividend price. This happens because trader m is in fact *not* "*marginal*" in their own market; because their portfolio constraint is binding they do not "price" asset m. Instead the local asset ends up being priced by the traders operating in other markets $n \neq m$. But these other traders do not care about the idiosyncratic risk of market m. Since assets are symmetric, they end up with the same equilibrium price.

6. Cross-sectional volatilities

We now pursue the implications of the general model with market-specific λ_m and hence a non-degenerate *cross-section of volatility*. Specifically, consider a finite number of market *types*. Each market contains the same number of assets, but there is a total measure ω_m of traders in market m with a supply *per trader* normalized to 1. With this notation, the aggregate endowment is $y = \sum_m y_m \omega_m$.

6.1. Calibration strategy and results

In the single λ benchmark, the value of λ was identified by matching a measure of systematic volatility, the return volatility of a well-diversified portfolio of stocks. Now a *vector* of segmentation parameters needs to be identified and this is achieved using a closely related strategy. In particular, market types are identified with quintile portfolios of stocks sorted on measures of idiosyncratic volatility. The value of λ_m for $m = 1, \ldots, 5$ is chosen to match the total volatility of the mth quintile *portfolio* as calculated by Ang et al. (2006). Similarly, the values of ω_m are chosen so that the average portfolio weight of the family in assets of market m matches the average market share for the mth quintile portfolio. Our procedure chooses these parameters simultaneously with the parameters of the stochastic process for cross-sectional endowment volatility. The values of the preference parameters and the aggregate growth parameters are kept at their benchmark values.

The calibrated parameters from this procedure are listed in Table 4. Market 1, with the lowest idiosyncratic volatility, has a segmentation parameter of only $\lambda_1 = 0.01$. This market consists of 20% of assets by number but it accounts for 51% of total market value. By contrast, market 5 has segmentation parameter $\lambda_5 = 0.37$ but accounts for only 2% of total market value. Across markets the segmentation parameters λ_m are monotonically increasing in m while the weights ω_m are monotonically decreasing in m. Averaging over the five markets $\overline{\lambda} = \sum_m \lambda_m \omega_m = 0.115$. Thus this economy, which matches the same aggregate moments as the benchmark model, hits its targets with an average amount of segmentation $\overline{\lambda} = 0.115$ roughly one-third that of the single parameter benchmark $\lambda = 0.31$. This suggests that there may be a significant bias when aggregating a collection of heterogeneously segmented markets into a "representative" segmented market.

6.2. Asset pricing implications

Table 5 shows the risk premia for each market type in the model and their empirical counterparts. In the data, the premium for the low volatility market 1 is 0.53% monthly (roughly 6.5% annual) whereas in the model it is 0.17% monthly. For markets with higher volatility, the model predicts that risk premia *monotonically* increase, reaching 0.53% monthly for market 5. However, the data exhibits a *hump-shaped* pattern for the cross-section of premia, reaching a maximum at about 0.69% monthly for market 3, then falling to -0.53% for the most volatile market 5. Thus the model does not account for the negative risk premia of the smallest, highest idiosyncratic volatility, markets.

Table 5 also shows the aggregate asset pricing implications of the model with market-specific λ_m. The aggregate equity premium is 2.9%, about 0.50% *higher* than in the benchmark single λ model, despite the fact that the average segmentation here is only $\overline{\lambda} = 0.115$, one-third the single λ benchmark. For comparison, the table shows the asset pricing implications for an otherwise identical single λ economy with $\lambda = \overline{\lambda} = 0.115$. The aggregation of the micro-segmentation frictions across the different markets adds some 1.2% annual to the equity premium, taking it from 1.7% to 2.9%.

Table 4
Market-specific segmentation: parameters and fit.

Panel A: *Segmentation parameters*

	Parameter		Moment			
			Portfolio std dev		Market share	
Market m	λ_m	ω_m	Data	Model	Data	Model
1	0.010	0.514	3.83	4.18	0.535	0.538
2	0.178	0.277	4.74	4.52	0.274	0.272
3	0.264	0.128	5.85	5.72	0.119	0.118
4	0.324	0.058	7.13	7.02	0.052	0.052
5	0.365	0.023	8.16	8.08	0.020	0.020
Average	0.115					

Panel B: *Idiosyncratic endowment volatility*

Parameter		Moment	Data	Model
$\overline{\sigma}$	0.816	Average cross-section std dev returns	16.40	16.46
$\sigma_{\epsilon v}$	0.198	Time-series std dev cross-section std dev returns	4.17	4.18
ϕ	0.891	AR(1) cross-section std dev returns	0.84	0.85

The top panel shows the five segmentation parameters λ_m and measures of traders ω_m, for $m=1,\ldots 5$, and the portfolio standard deviation and market share moments in Ang et al. (2006) data they are chosen to match. The bottom panel shows the idiosyncratic endowment volatility process parameters and the moments in Goyal and Santa-Clara (2003) cross-sectional standard deviation of stock returns data they are chosen to match.

Table 5
Asset pricing implications of market-specific segmentation.

Panel A: *Market-specific asset pricing implications*

	Risk premia	
Market m	Data	Model
1	0.53	0.17
2	0.65	0.23
3	0.69	0.33
4	0.36	0.44
5	−0.53	0.53

Panel B: *Aggregate asset pricing implications*

			Model	
	Moment	Data	λ_m	$\overline{\lambda}$
Equity premium	$\mathbb{E}[R_M-R_f]$	5.27	2.92	1.69
	$\text{Std}[R_M-R_f]$	14.25	15.54	11.11
Sharpe ratio	$\mathbb{E}[R_M-R_f]/\text{Std}[R_M-R_f]$	0.38	0.17	0.14
Market return	$\mathbb{E}[R_M]$	7.24	11.07	10.27
	$\text{Std}[R_M]$	14.44	16.16	11.56
Risk free rate	$\mathbb{E}[R_f]$	1.81	8.15	8.58
	$\text{Std}[R_f]$	1.20	3.65	2.92
Price/dividend ratio	$\mathbb{E}[p/y]$ (annual)	34.38	13.90	14.04
	$\text{Std}[\log(p/y)]$ (annual)	38.63	32.84	23.22
	$\text{Auto}[\log(p/y)]$ (monthly)	0.99	0.88	0.88

The top panel shows the market risk premia implied by the five markets $m=1,\ldots 5$ and their counterparts in Ang et al. (2006) data. These are reported as monthly percent. The bottom panel shows the aggregate asset pricing implications. The column marked λ_m refers to the model with market-specific segmentation parameters while the column marked $\overline{\lambda}$ refers to a model with a single segmentation parameter λ that is set equal to the mean $\overline{\lambda}=\sum_m \lambda_m \omega_m$ of the market-specific λ_m model.

6.3. *Welfare costs of market segmentation*

We measure the welfare costs of segmentation as the percentage increase in lifetime consumption required to make the family indifferent between living with a given amount of market segmentation or eliminating that segmentation entirely (the same way that Lucas, 1987, measures the welfare costs of business cycles).

For the single λ benchmark the welfare cost of segmentation is 2.2% of lifetime consumption. Fluctuations in σ_t account for a small yet economically significant share of this cost. If the calculation is repeated with $\sigma_t = \overline{\sigma}$, then the cost of segmentation drops to 1.85% of lifetime consumption. Fig. 5 shows the cost of segmentation as a function of the segmentation parameter λ. The cost of segmentation is increasing and *convex* in λ; traders find it increasingly costly to bear more idiosyncratic volatility. This suggests that, in a multiple asset model, the average level of segmentation is likely to underestimate the true economic cost of segmentation.

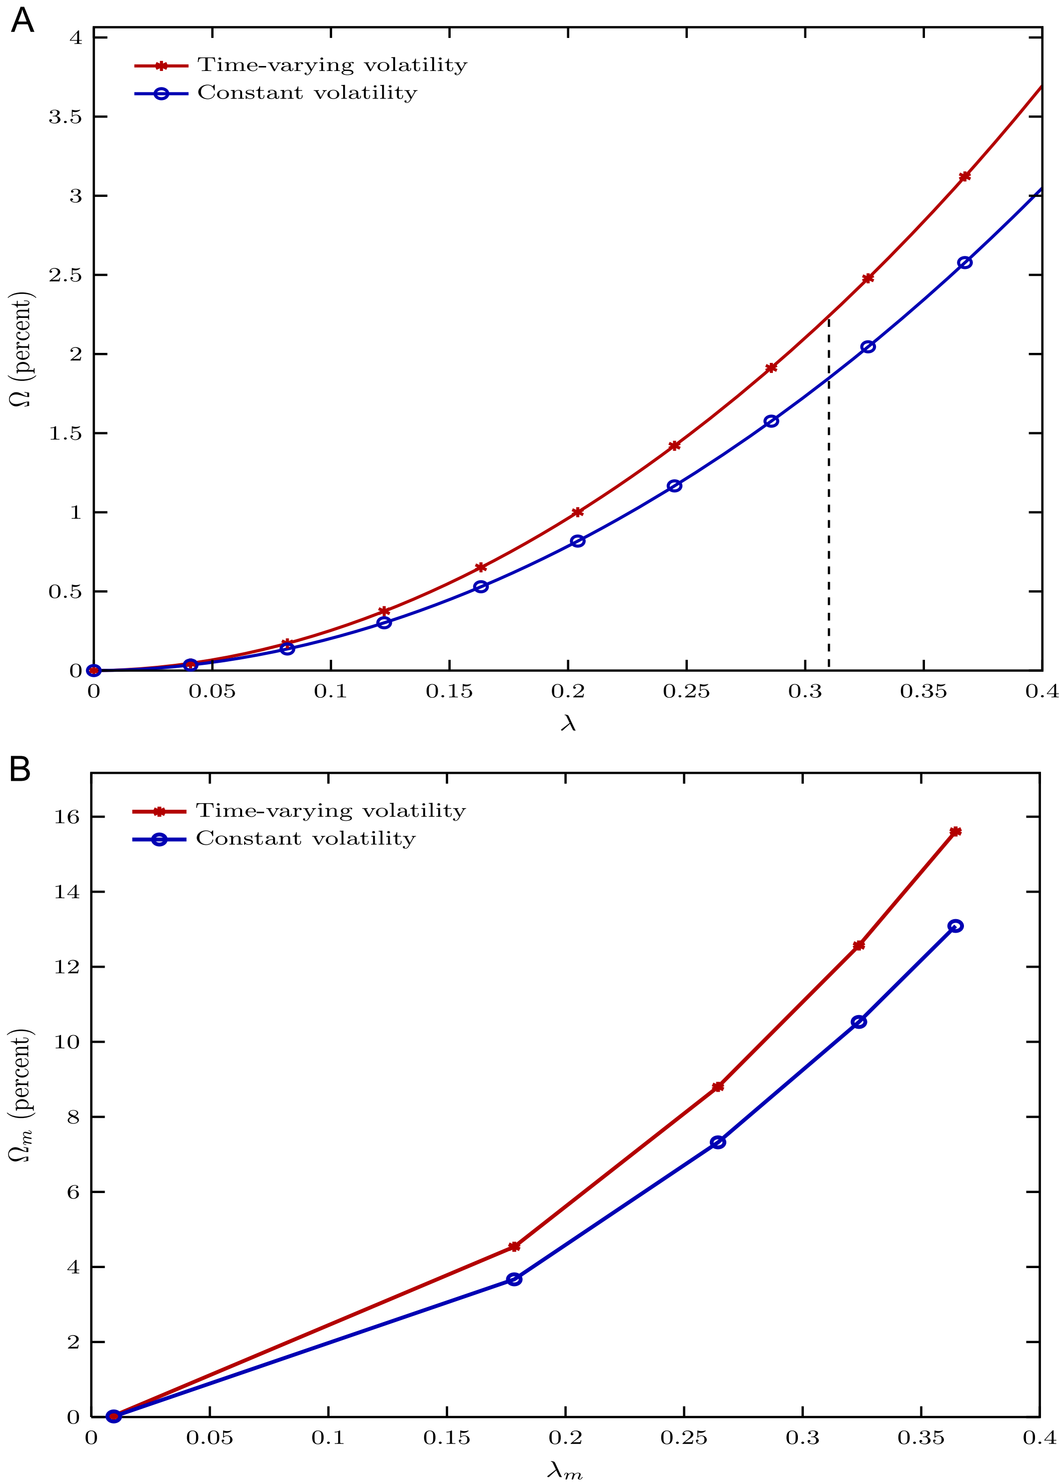

Fig. 5. Welfare costs of segmentation. Panel A: *Single segmentation parameter* λ. For our benchmark calibration with $\lambda = 0.31$, the welfare cost of segmentation is 2.2% of lifetime consumption. Time-varying volatility accounts for a small but economically significant share of this cost. If σ_t is constant, the cost of segmentation drops to 1.85% of lifetime consumption. Panel B: *Market-specific segmentation parameters* λ_m. The welfare cost of segmentation in each market with and without time-varying volatility. The average welfare cost of segmentation is 3% of lifetime consumption, higher than in the single $\lambda = 0.31$ benchmark, despite the average segmentation $\overline{\lambda} = 0.11$ being only one-third as high.

Indeed, for the multiple λ_m economy the welfare cost is 3% of lifetime consumption, considerably larger than the 2.2% for the single λ economy. The welfare cost is higher than in the single λ case because of two effects. First, the level of volatility is larger in the multiple λ_m calibration than in the single λ calibration. This increases the cost of segmentation for any λ. Second, the cost is a convex function of λ_m so that Jensen's inequality implies that increased dispersion in segmentation raises the welfare cost. Appendix F provides more detail on these calculations.

7. Conclusion

To assess the aggregate implications of market-specific frictions, we develop a consumption-based asset pricing model in which assets are traded in a financial system consisting of many segmented markets. Because of the segmentation, a trader operating in one particular market cannot fully diversify the idiosyncratic risk specific to that market. Assets in each micro market are priced by a convex combination of the individual marginal utility of traders specialized in that asset (who bear some idiosyncratic risk), and the average marginal utility in the economy (reflecting diversification of the remaining idiosyncratic risk in a large portfolio).

Our model implies that market-specific segmentation frictions can have significant implications for aggregate asset prices. The amount of segmentation is calibrated to reproduce key facts on systematic and idiosyncratic return volatility and the model then implies a sizeable aggregate equity premium and pronounced time-variation in expected returns. A disaggregated version of the model that allows the amount of segmentation to differ across markets produces the same aggregate asset pricing implications but with a much smaller average amount of segmentation. Moreover, despite having a smaller average amount of segmentation, this disaggregated version of the model also implies a significantly larger welfare cost of segmentation. In short, both the mean and the cross-sectional dispersion of the segmentation friction matter in the aggregate.

Finally, our segmented markets model has markedly different asset pricing implications from those of an otherwise similar incomplete markets model. Idiosyncratic risk leads to a significant aggregate risk premium in the segmented markets model but has no such implications in the incomplete markets model.

Acknowledgments

We thank Andrew Atkeson, David Backus, Bruno Biais, Alexandre Dmitriev, Xavier Gabaix, Johan Hombert, Greg Kaplan, Stijn Van Nieuwerburgh, Gianluca Violante and seminar participants at the ANU, FRB Richmond, FRB Philadelphia, NYU, Ohio State, Sydney, UNSW, UTS, Wharton, Toulouse School of Economics, the LAEF UCSB conference on Financial Frictions and Segmented Asset Markets, the 2011 Southern Workshop in Macroeconomics, the 2010 Australian Conference on Quantitative Macroeconomics, the 2009 Australasian Macroeconomics Workshop, the 2009 Sydney–Melbourne Workshop on Macroeconomic Theory, and the 2007 and 2008 SED annual meetings for helpful comments and conversations. We especially thank Urban Jermann and an anonymous referee whose suggestions greatly improved the paper and Amit Goyal and Turan Bali for sharing their data with us. Chris Edmond gratefully acknowledges support from the Australian Research Council, Grant DP-110103457. Pierre-Olivier Weill gratefully acknowledges support from the National Science Foundation, Grant SES-0922338.

Appendix A. Supplementary material

Supplementary data associated with this article can be found in the online version at http://dx.doi.org.10.1016/j.jmoneco.2012.03.006.

References

Aiyagari, S.R., Gertler, M., 1991. Asset returns with transaction costs and uninsurable individual risks. Journal of Monetary Economics 27, 309–331.
Aiyagari, A.R., Gertler, M., 1999. Overreaction of asset prices in general equilibrium. Review of Economic Dynamics 2, 3–35.
Alvarez, F., Atkeson, A., Kehoe, P.J., 2002. Money, interest rates, and exchange rates with endogenously segmented markets. Journal of Political Economy 110, 73–112.
Ang, A., Hodrick, R.J., Xing, Y., Zhang, X., 2006. The cross-section of volatility and expected returns. Journal of Finance 61, 259–299.
Backus, D.K., Zin, S.E., 1994. Reverse Engineering the Yield Curve. NBER Working Paper 4676.
Bali, T.N., Cakici, X., Yan, S., Zhang, Z., 2005. Does idiosyncratic risk really matter? Journal of Finance 60, 905–929.
Basak, S., Cuoco, D., 1998. An equilibrium model with restricted stock market participation. Review of Financial Studies 11, 309–341.
Breeden, D., 1979. An intertemporal asset pricing model with stochastic consumption and investment opportunities. Journal of Financial Economics 7, 265–296.
Campbell, J.Y., Cochrane, J.H., 1999. By force of habit: a consumption-based explanation of aggregate stock market behavior. Journal of Political Economy 107, 205–251.
Campbell, J.Y., Lettau, M., Malkiel, B.G., Xu, Y., 2001. Have individual stocks become more volatile? An empirical exploration of idiosyncratic risk. Journal of Finance 56, 1–43.
Chien, Y.-L., Lustig, H., 2010. The market price of aggregate risk and the wealth distribution. Review of Financial Studies 23, 1596–1650.

Chien, Y.-L., Cole, H., Lustig, H., 2011. A multiplier approach to understanding the macro implications of household finance. Review of Economic Studies 78, 199–234.
Collin-Dufresne, P., Goldstein, R.S., Martin, J.S., 2001. The determinants of credit spread changes. Journal of Finance 56, 2177–2207.
Edmond, C., Weill, P.-O., 2012. Aggregate implications of micro asset market segmentation. Online supplementary appendix.
Gabaix, X., Krishnamurthy, A., Vigneron, O., 2007. Limits of arbitrage: theory and evidence from the mortgage-backed securities market. Journal of Finance 62, 557–595.
Gomes, F., Michaelides, A., 2008. Asset pricing with limited risk sharing and heterogeneous agents. Review of Financial Studies 21, 415–448.
Goyal, A., Santa-Clara, P., 2003. Idiosyncratic risk matters! Journal of Finance 58, 975–1007.
Guvenen, F., 2009. A parsimonious macroeconomic model for asset pricing. Econometrica 77, 1711–1754.
He, H., Modest, D.M., 1995. Market frictions and consumption-based asset pricing. Journal of Political Economy 103, 94–117.
He, Z., Krishnamurthy, A., 2010a. Intermediary Asset Pricing, Working Paper.
He, Z., Krishnamurthy, A., 2010b. A Model of Capital and Crises, Working Paper.
Heaton, J., Lucas, D.J., 1996. Evaluating the effects of incomplete markets on risk sharing and asset pricing. Journal of Political Economy 104, 443–487.
Kocherlakota, N.R., 1990. On the 'discount' factor in growth economies. Journal of Monetary Economics 25, 43–47.
Kocherlakota, N.R., Pistaferri, L., 2009. Asset pricing implications of Pareto optimality with private information. Journal of Political Economy 117, 555–590.
Krueger, D., Lustig, H., 2010. When is market incompleteness irrelevant for the price of aggregate risk (and when is it not)? Journal of Economic Theory 145, 1–41.
Kyle, A.S., Xiong, W., 2001. Contagion as a wealth effect. Journal of Finance 56, 1401–1440.
Lucas Jr., R.E., 1978. Asset prices in an exchange economy. Econometrica 46, 1429–1445.
Lucas Jr., R.E., 1987. Models of Business Cycles. Blackwell, Cambridge.
Lucas Jr., R.E., 1990. Liquidity and interest rates. Journal of Economic Theory 50, 237–264.
Lustig, H., Van Nieuwerburgh, S., 2005. Housing collateral, consumption insurance and risk premia: an empirical perspective. Journal of Finance 60, 1167–1219.
Luttmer, E.G., 1996. Asset pricing in economies with frictions. Econometrica 64, 1439–1467.
Luttmer, E.G., 1999. What level of fixed costs can reconcile consumption and stock returns? Journal of Political Economy 107, 969–997.
Mankiw, N.G., Zeldes, S.P., 1991. The consumption of stockholders and nonstockholders. Journal of Financial Economics 29, 97–112.
Mehra, R., Prescott, E.C., 1985. The equity premium: a puzzle. Journal of Monetary Economics 15, 145–161.
Pavlova, A., Rigobon, R., 2008. The role of portfolio constraints in the international propagation of shocks. Review of Economic Studies 75, 1215–1256.
Piazzesi, M., Schneider, M., Tuzel, S., 2007. Housing, consumption and asset pricing. Journal of Financial Economics 83, 531–569.
Ravn, M.O., Uhlig, H., 2002. On adjusting the Hodrick–Prescott filter for the frequency of observations. Review of Economics and Statistics 84, 371–375.
Saito, M., 1995. Limited Participation and Asset Pricing. University of British Columbia, Working Paper.
Shleifer, A., Vishny, R.W., 1997. The limits of arbitrage. Journal of Finance 52, 35–55.
Storesletten, K., Telmer, C.I., Yaron, A., 2004. Cyclical dynamics in idiosyncratic labor market risk. Journal of Political Economy 112, 695–717.
Tauchen, G., Hussey, R., 1991. Quadrature based methods for obtaining approximate solutions to nonlinear asset pricing models. Econometrica 59, 371–396.
Vayanos, D., 2005. Flight to Quality, Flight to Liquidity, and the Pricing of Risk. London School of Economics, Working Paper.
Weil, P., 1989. The equity premium puzzle and the risk-free rate puzzle. Journal of Monetary Economics 24, 401–421.