Structural Change and the Long-run Dynamics of the Equity Premium

Vasco Carvalho* Vincenzo Merella† Alessio Moro‡ Loris Rubini§

VERY PRELIMINARY DRAFT. DO NOT QUOTE.

Abstract

We study whether structural change helps in rationalizing the declining equity premium observed in the data in the post-war period in the U.S. The recent literature on micro shocks and aggregate fluctuations finds that, when the portfolio of production technologies in use in an economy evolves over time, so do aggregate fluctuations. In particular, technological diversification tracks well low frequency movements in aggregate and industry level U.S. asset prices. We investigate whether the relation between structural change in aggregate consumption and aggregate volatility has implications for risky asset returns and their differentials with the risk-free return. Our preliminary results suggest that a generalized multi-sector Lucas-tree model, which can account for the structural change observed in the U.S., is able to pin down the low frequency dynamics of the U.S. equity premium over the period 1947-2010.

JEL Classification: C68, G12, O41

Keywords: Asset Pricing, Equity Premium, Structural Change.

*University of Cambridge. E-mail: vmcarvalho.web@gmail.com.
†University of Cagliari. E-mail: merella@unica.it.
‡University of Cagliari. E-mail: amoro@unica.it.
§University of New Hampshire. E-mail: loris.rubini@unh.edu.

1 Introduction

A direct implication of the recent literature on micro shocks and aggregate fluctuations is that, if one is willing to accept the view that the micro-structure of production has a bearing on business cycle properties, it should also be the case that, as this structure evolves over time, so does the nature of aggregate risk and fluctuations. In particular, it is well known that over the last half-century advanced economies have witnessed considerable changes in their technology portfolio, with the decline of the goods sector and the rise of services.[1]

In this paper, we aim to study whether the secular process of structural change observed in modern economies helps to rationalize the low frequency dynamics of the equity premium as observed in data, thus providing a link between the literature on micro shocks and aggregate fluctuations in a multi-sector settings and the emerging literature on technological change and asset prices (e.g. Garleanu et al. (2012), or Papanikolaou (2011)).[2]

Our analysis is grounded on two facts identified for the U.S. economy in the literatures of asset pricing and structural change, respectively. On the one hand, the equity premium fell considerably in the post-war period and, while it may have risen modestly in the early twenty-first century (Campbell (2008)), it continues to hover at its lowest levels since the beginning of the twentieth century (Jordà et al. (2019)). On the other hand, the share of services in aggregate consumption (as well as in GDP) increased substantially over the post-war period (Herrendorf et al. (2014)).

A link between these two facts is suggested by a number of theoretical and quantitative findings in the two literatures. First, structural change between goods and services induces a decline in both aggregate GDP volatility (Moro (2012), and Moro (2015)) and the real interest rate (Leon-Ledesma and Moro (2020)) along the growth path. Second, aggregate volatility is heightened during periods of low diversification of production, i.e. when a sector accounts for a disproportionate share of production (Carvalho and Gabaix (2013)).[3] Third, the link between the evolution of aggregate volatility over time and the equity premium is consistent with the findings in Lettau et al. (2008), who have shown that macroeconomic risk, as measured by the conditional volatility of GDP, and aggregate valuation ratios are strongly correlated.

Our main conjecture from the above results is that, as structural change influences aggregate volatility —which in turn affects both the risky return and the risk-free interest rate— it might also influence the long run dynamics of the equity premium. In this paper

[1] See Herrendorf et al. (2014) for a summary.

[2] See Kogan and Papanikolaou (2012) for a review.

[3] See also Imbs and Wacziarg (2003), Koren and Tenreyro (2007) and Koren and Tenreyro (2013) for concepts and consequences of 'technological diversification'.

we explore this link formally by augmenting the workhorse Lucas-tree asset pricing model to an environment with two non-stationary trees (sectors, which we label *goods* and *services*), whose size and fruits (dividends) are determined by the laws of motion of structural change growth models. Thus, we build on the contributions of Cochrane et al. (2008) and Martin (2013) and extend them to settings where the low frequency dynamics of asset prices—and returns— are the endogenous result of a well specified quantitative structural change model.

Formally, agents receive in each period their income from sector-specific, idiosyncratic dividend payments, and solve the problem of maximizing their expected lifetime utility by choosing how much of their income is allocated to consumption and sectoral asset shares. Specifically, we explore a setting with non-homothetic preferences and heterogeneous sectoral endowment processes, which the literature on structural change has identified as key drivers of long-term sectoral dynamics. Our setting provides for potential multiple mechanisms affecting equity premium dynamics: not only, as discussed above, does the endogenous evolution of technologies induce non-stationary dynamics for aggregate risk and the risk free rate; but also the interaction of sectoral growth with non-homothetic preferences generates low frequency changes in agents' risk aversion.[4]

More specifically we consider a type of preference structure *à la* Boppart (2014) that can account for long run structural change in the economy. In this setting, preferences imply that the two goods, and so the fruits of the two trees, are imperfect substitutes. This implies that the relative quantity of fruits in the two sectors determines the relative price of the two consumption goods which, in turn, affects the prices of the two types of assets (by influencing the value of dividend value). Thus, with respect to a one-sector model, or a multi-sector model with perfectly substitutable fruits, in our setting the relative abundance of fruits interact with their prices and, as a result, on sectoral asset (i.e. trees) prices, generating non-obvious equity premium dynamics.

Based on this setup, we aim at a quantitative exploration of the link between structural change and asset prices. In particular, we investigate whether, by calibrating the model to long-run structural change dynamics in the U.S., we can simultaneously account for both the low frequency movements in the equity premium and the long term evolution of sector shares in consumption. Technically, the quantitative exercise is not trivial since, in a stochastic asset pricing environment like the one we employ, structural change implies that the agents' expectations must take into account the prospective changes in the structure of the economy. As a result, the time-independent recursive structure of the Lucas-tree model's solution cannot be obtained. In order to resolve this issue we resort to a new parametrized expectations

[4]See Ait-Sahalia et al. (2004), for an early exploration of the role of non-homothetic preferences in accounting for the equity premium.

algorithm applied to stochastic structural change environments, as recently developed by Rubini and Moro (2019).

We calibrate the model using U.S. data and compare them with the empirical evidence that we construct on the equity premium following the methodology in (Jordà et al. (2019)). Our empirical analysis confirms that there has been a marked decline of the EP in U.S. data in the post-war, with its current value being roughly 50% of the initial value. The evolution of the equity premium in our setting confirms this broad pattern by exhibiting a decline at the early stages of the transition from the goods sector to the services sector and a stabilisation afterwards. Quantitatively, the equity premium declines from 4.5% to 2% in the model over the sample period, in line with what we observe in the data.

The reminder of the paper is as follows. Section 2 presents the evidence on the equity premium in the U.S. and the world in the post-war period; Section 3 presents the theoretical framework; Section 4 presents the main results and Section 5 concludes.

2 Empirical Evidence

Empirical estimates of the equity risk premium typically take three steps:

1. Estimate the expected return on equity
2. Estimate the expected returns on a risk-free rate
3. Take the difference to calculate the equity risk premium.

In order to calculate expectations of the likely future returns on equity (Step 1), we use a total-returns-based approach to estimate the expected return on equity, applied to two different datasets for the U.S. stock market data . Regarding Step 2, recent studies take a long-run measure of TIPS as the safe U.S. interest rate, arguing that this market is the most liquid associated which will return a real yield. We follow the earlier methodologies, using the USTs directly given that we wish to document long-run historical series (TIPS were first issued in 1997).

At any point in time the exact calculation of the equity risk premium may then be given as:

$$\text{Equity Risk Premium}_t = \frac{p_t^{\text{Equity}} + d_t}{p_{t-1}^{\text{Equity}}} - \frac{p_t^{\text{Safe}} + i_t}{p_{t-1}^{\text{Safe}}},$$

where equity is assumed to pay a dividend and the safe assets is a (constant maturity) nominal treasury bond, paying interest.

In a final step we account for the sizeable fluctuations in these measures by following the literature to compute a rolling average. Campbell (2008) uses a 3-year rolling average to smooth the data whereas Jordà et al. (2019) use a 10-year average. We use a 10-year average as our primary data source is Jordà et al. (2019). Varying the estimation window makes little difference for the results below.

The majority of the data sources are taken directly from Jordà et al. (2019), who provide a rich description of the return on a number of asset classes across a number of advanced economies. Their data has been collected from a wide range of national sources. All data used are annual.

We construct the consumption shares of goods and services in total consumption expenditure by taking the ratio of their nominal expenditure over total consumption expenditure in each year in the U.S.[5]

Figure 1 reports our estimate for the equity premium and for the two shares.

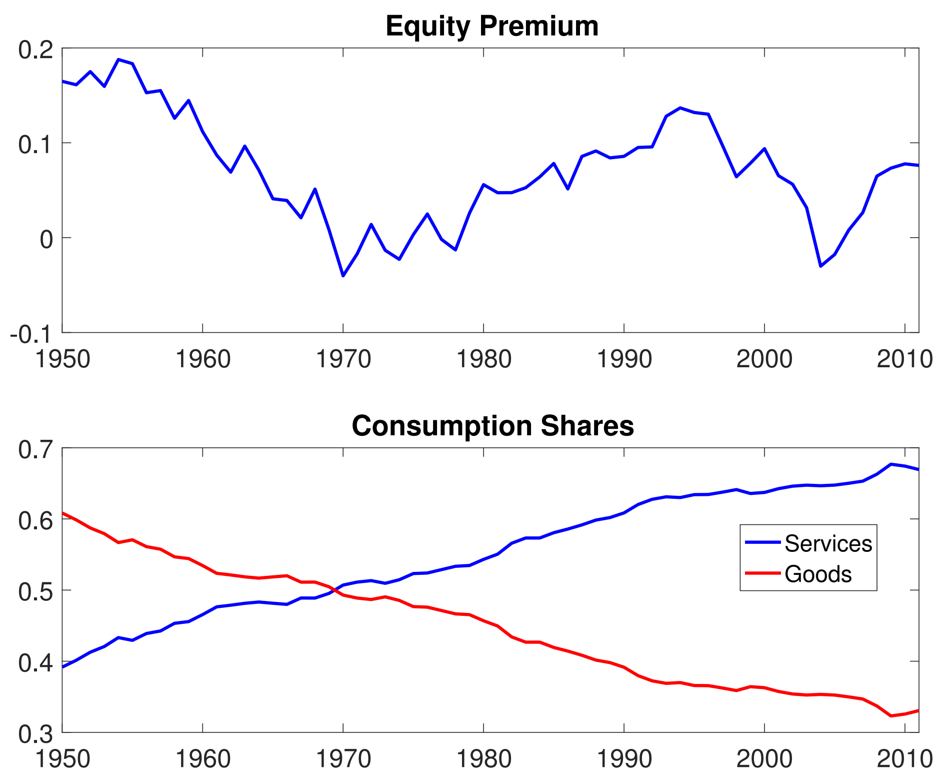

Figure 1: Equity Premium (top) and Consumption shares in the U.S. (bottom). Source: Jordà et al. (2019), PWT, BEA and own calculations.

[5]BEA table 1.1.5.

3 Theoretical Framework

We first describe the assets and the probability structure, together with the budget constraint of the consumer. Next, we describe the agents' preferences and the resulting Euler Equation.

3.1 Financial Structure

Time is indexed by t. There are two trees in the economy, labelled goods and services and indexed by $z = m, s$. Each tree yelds a distinct (sectoral) stochastic fruit (dividends) in quantity $y_{z,t}$ in each period. Thus, a tree represents both a specific asset type whose shares are traded in the market, but also a sector supplying a specific type of consumption good (the fruit), which is different from the consumption good produced by the other tree (i.e. fruits are not perfect substitutes). The growth factor of the sectoral dividend is

$$g_{z,t+1} \equiv y_{z,t+1}/y_{z,t} = g_z + \varepsilon_{z,t+1}, \ z = m, s \tag{1}$$

where g_m and g_s are the deterministic portions of the sectoral growth factors, potentially different across sectors and $\varepsilon_{m,t+1}$ and $\varepsilon_{s,t+1}$ are stochastic components.

The budget constraint of any agent at time t is

$$p_{m,t}c_{m,t} + p_{s,t}c_{s,t} = p_{m,t}y_{m,t}a_{m,t} + p_{s,t}y_{s,t}a_{s,t} - q_{m,t}\left(a_{m,t+1} - a_{m,t}\right) - q_{s,t}\left(a_{s,t+1} - a_{s,t}\right). \tag{2}$$

At each t the consumer holds a certain amount $a_{z,t}$ of the asset backed by tree z from which it receives the market value of dividends $p_{z,t}y_{z,t}$, changes its asset position by $(a_{z,t+1} - a_{z,t})$ at the market price of the asset $q_{z,t}$, and purchases fruits of each sector $c_{z,t}$ for consumption purposes at price $p_{z,t}$.

The asset return is defined as

$$R_{z,t+1} \equiv \frac{x_{z,t+1}}{q_{z,t}}. \tag{3}$$

We define R_t as the aggregate return, that is, the return obtained by holding a portfolio comprising $a_{m,t+1}$ m-sector assets and $a_{s,t+1}$ s-sector assets as

$$R_{t+1} \equiv \frac{x_{m,t+1}a_{m,t+1} + x_{s,t+1}a_{s,t+1}}{q_{m,t}a_{m,t+1} + q_{s,t}a_{s,t+1}}. \tag{4}$$

This is computed as the sum of the product of payoffs and holdings of the two assets at time $t + 1$, divided by the cost of those assets at time t.

In equilibrium, the representative agent holds the entire amount of assets in the economy. As the total amount of each asset is normalized to one in the economy, we can define the *representative agent's sectoral portfolio weights* as

$$\phi_t \equiv \frac{q_{m,t}}{q_{m,t}+q_{s,t}}; \quad 1-\phi_t = \frac{q_{s,t}}{q_{m,t}+q_{s,t}}.$$

Further defining the relative asset price as

$$q_t \equiv \frac{q_{m,t}}{q_{s,t}},$$

we can redefine sectoral portfolio weights as

$$\phi_t \equiv \frac{q_t}{1+q_t}; \quad 1-\phi_t = \frac{1}{1+q_t}. \tag{5}$$

This allows us to re-write the aggregate returns of the portfolio as a weighted average of sectoral risky returns as

$$R_{t+1} = \phi_t R_{m,t+1} + (1-\phi_t) R_{s,t+1}.$$

By taking expectations at time t yields

$$E_t\{R_{t+1}\} \equiv \phi_t E_t\{R_{m,t+1}\} + (1-\phi_t) E_t\{R_{m,t+1}\}. \tag{6}$$

By defining the risk-free return as R^f_{t+1} we can define the sectoral equity premia as

$$EP_{z,t+1} \equiv E_t\{R_{z,t+1}\} - R^f_{t+1}, \ z = m, s \tag{7}$$

and the aggregate equity premium as

$$EP_{t+1} \equiv E_t\{R_{t+1}\} - R^f_{t+1} = \phi_t EP_{m,t+1} + (1-\phi_t) EP_{s,t+1}. \tag{8}$$

3.2 Preferences

We assume that therepresentative agent's preferences are represented by a one-period indirect utility function as in Boppart (2014):

$$V(p_{s,t}, p_{m,t}, E_t) = \frac{1}{\epsilon}\left[\frac{E_t}{p_{s,t}}\right]^{\epsilon} - \frac{\nu}{\gamma}\left(\frac{p_{m,t}}{p_{s,t}}\right)^{-\gamma} - \frac{1}{\epsilon} + \frac{\nu}{\epsilon}, \tag{9}$$

where E_t is consumption expenditure given by

$$p_{m,t}c_{m,t} + p_{s,t}c_{s,t} = E_t,$$

and parameter restrictions are $0 \leq \epsilon \leq \gamma \leq 1$ and $\nu > 0$.

Boppart (2014) and Leon-Ledesma and Moro (2020) show that introducing these preferences in a growth environment in which the relative price $p_{m,t}/p_{s,t}$ deterministically grows at a constant rate as in U.S. data (due to different TFP growth in the two sectors), allows the model to match both 1) the nominal share of services in consumption expenditure and 2) the ratio of real services consumption to goods consumption. That is, in an environment in which the relative price is exogenosly matched, these preferences allow to match the evolution of the relative quantity (and so the evolution of the nominal shares). In our framework, as in a typical asset pricing model, quantities are exogenously given through trees fruits, and consumption prices are endogenously determined. Consumption prices in turn are one of the determinants of asset prices in this setting. For this reason, it is of crucial importance to introduce a preference specification that can account for the long-run evolution of consumption prices given the exogenous evolution of consumption.

The demand functions implied by the preference representation in (9) are given by

$$c_s(p_{s,t}, p_{m,t}, E_t) = \frac{E_t}{p_{s,t}} - \nu E_t^{1-\epsilon} p_{m,t}^{\gamma-\epsilon} p_{s,t}^{\epsilon-\gamma-1} \tag{10}$$

$$c_m(p_{s,t}, p_{m,t}, E_t) = \nu E_t^{1-\epsilon} p_{m,t}^{\gamma-\epsilon-1} p_{s,t}^{\epsilon-\gamma} \tag{11}$$

and market clearing implies

$$c_z(p_{m,t}, p_{s,t}, E_t) = y_{z,t},\ z = m, s.$$

The intertemporal problem of the representative agent is to maximize the discounted sum of one-period utility functions

$$\sum_{t=0}^{\infty} \beta^t V(p_{m,t}, p_{s,t}, E_t),$$

subject to the budget constraint (2). From the first order conditions of this problem, we obtain the price of each asset $z = m, s$, given by

$$q_{z,t} = \mathbb{E}\left(\frac{\beta V_E(p_{m,t+1}, p_{s,t+1}, E_{t+1})}{V_E(p_{m,t}, p_{s,t}, E_t)} (p_{z,t+1} y_{z,t+1} + q_{z,t+1})\right). \tag{12}$$

The problem with solving the model is that the expectation in the right hand side of equation (12) is not known. In a typical asset pricing model, the value of $q_{z,t}$ does not depend on t, as the problem is either stationary or can be easily made stationary. Here, instead, the value of $q_{z,t}$ depends on t, because of the presence of structural change. Loosely speaking, the expression inside the expectation term (12) is non-stationary because of the presence of structural change in the model. As a result, the expectation on the right hand side is different at each t, and so is the price of the asset $q_{z,t}$. For this reason, in the next section we resort to the parametrized expecations algorithm developed in Rubini and Moro (2019) to solve the stochastic structural change model.

4 Quantitative Analysis

4.1 Data

We briefly introduce the set of U.S. data that we use to calibrate the model. Sectoral and aggregate consumption levels are drawn from the National Income and Product Accounts (NIPA) database, released by the Bureau of Economic Analysis (BEA). Specifically, we use quarterly data on Personal Consumption Expenditures on (i) Aggregate, on (ii) Goods and on (iii) Services (Table 1.1.5, rows 2, 3 and 6, respectively) from 1947:Q1 to 2010:Q2 to obtain nominal consumption values; and quarterly data on Quantity Indexes of Personal Consumption Expenditures on (iv) Aggregate, on (v) Goods and on (vi) Services (Table 1.1.3, again rows 2, 3 and 6, respectively) over the same period to obtain the real consumption values. We compute the nominal sectoral shares as the ratio of (ii) over (i) for (a) goods and of (ii) over (i) for (b) services; likewise, we obtain real sectoral share as the ratios of (v) over (vi) for (c) goods and (vi) over (iv) for (d) services. We then compute the relative (1) nominal and real (2) shares of services over goods as the ratio of (b) over (a) and (d) over (c), respectively. Finally, we compute the relative price of services over goods as the ratio of (1) over (2).

4.2 Transition Probability Matrix

As we mentioned at the beginning of Section 3, the stochastic behavior of the model is driven by $\varepsilon_{z,t+1}$, with $z = m, s$, which is a binary random component defined by

$$\varepsilon_{z,t+1} = \begin{cases} \varepsilon_z & (+) \\ -\varepsilon_z & (-) \end{cases}, \quad z = m, s \tag{13}$$

Since each stochastic process is binary, we have four possible joint outcomes

$$\begin{array}{ll} (+,+) & \text{state 1} \\ (+,-) & \text{state 2} \\ (-,+) & \text{state 3} \\ (-,-) & \text{state 4} \end{array}$$

where the first element refers to sector m, the second to sector s. We assume that a Markovian process governs the probability distribution associated to these four states. Future outcomes thus depend on the current state: we have the set of *conditional* probabilities, which we collect in a (4×4) matrix $\{\pi_{ij}\}$, with $i, j = \{1, 2, 3, 4\}$, where each element depicts the likelihood that the economy will be in the particular future state j, given that the current state is i; and a set of unconditional probabilities, which we collect in the (1×4) vector $\{\pi_j\}$.

4.3 Calibration

We calibrate the model as follows. One model period represents a quarter. We use data starting in the first quarter of 1950 and ending in the second quarter of 2018. Following the asset pricing literature, we set $\beta = 0.98$.. The unconditional growth rates of each tree are set to replicate the real growth in consumption of goods and consumption of services during the sameperiod in the U.S. In the data, consumption in thegoods sector grows at an annual rate of 1.7%, and in the service sector at 2.1%. Thus, we set the quarterly growth rates to $g_m = 0.42\%$ and $g_s = 0.52\%$. The remaining parameters in the technology side relate to the stochastic shocks, which include their magnitude and the transition probabilities.

The transition probabilities are calibrated as follows. First, we establish the trends of sectoral consumption levels, separating them from cyclical components via Hodrick-Prescott filtering applied to the time series of real consumption of goods and of services, respectively. Second, we compute the deviations of the actual sectoral consumption levels from the relevant trend. Third, we identify the realized state of nature at each point in time by assigning to it the element of the set $\{1, 2, 3, 4\}$ corresponding to the pair of signs of the relevant sectoral consumption deviations. Fourth, we compute the frequency of each state occurring in the dataset, *conditional* on the lagged realized state. We thus obtain a (4×4) matrix of frequencies, which we convert to the transition probability matrix $\{\pi_{ij}\}$ by dividing each element by the relevant number of total occurrences of the lagged realized state. Finally, we obtain the unconditional probability vector $\{\pi_j\}$ by applying a sufficiently large power to $\{\pi_{ij}\}$ so that the conditional probability matrix converges to a set of four identical vectors

Parameter	Target	Value
β	Standard Value	0.9800
g_m	Annual growth rate in goods of 1.7%	0.0042
g_s	Annual growth rate in services of 2.1%	0.0052
ε_m	Std. dev. of aggregate consumption of 0.012	0.0173
ε_s	Std. dev. of goods relative to services of 2.32	0.0065
γ	Leon-Ledesma and Moro (2020)	0.5053
ϵ	Leon-Ledesma and Moro (2020)	0.2179
ν	Leon-Ledesma and Moro (2020)	0.6399

Table 1: Calibrated Parameters and Targets

—each in fact depicting $\{\pi_j\}$. This results in the following transition matrix:

$$\Pi = \begin{bmatrix} 0.360 & 0.220 & 0.290 & 0.130 \\ 0.230 & 0.190 & 0.100 & 0.480 \\ 0.305 & 0.195 & 0.240 & 0.260 \\ 0.210 & 0.210 & 0.130 & 0.450 \end{bmatrix}$$

Each element in row i and column j represents the probability that next period's state is j given than today's state is i. Given these probabilities, we set the magnitude of the shocks to match the average aggregate volatility of consumption and the volatility of services relative to goods in the period considered. To compute these volatilities, we obtain quarterly chain weighted consumption expenditures from the BEA (Table 1.1.6), take logarithms, and detrend using an HP filter with smoothing parameter 1,600. The resulting standard deviation of consumption is 0.012. Similarly, we obtain data on the consumption of goods and services from the BEA Table 1.2.6. Following the same procedure, we compute the standard deviation of goods to be 2.32 times the standard deviation of services. We match these numbers by setting $\varepsilon_m = 0.0173$ and $\varepsilon_s = 0.0065$.

The utility function parameters are taken from Leon-Ledesma and Moro (2020), who set them to match the consumption shares in 1950 and in 2010 by feeding the model with the observed growth in the relative price services/goods. The values are $\gamma = 0.5053, \epsilon = 0.2179$ and $\nu = 0.6399$. Table 1 summarizes the calibration strategy.

4.4 Results

Our conjecture of a link between structural change and asset prices stems from the following considerations. On the one hand, when the technology portfolio is poorly diversified and concentrated in a narrow range of sectors, shocks to these sectors disproportionally affect GDP, which becomes more responsive to idiosyncratic fluctuations. On the other hand, when output is distributed among more sectors (high diversification), sectoral shocks tend to offset each other. As a result, as structural transformation influences both aggregate volatility (which affects the risky return) and the risk-free interest rate, our conjecture is that it might also influence the equity premium.

Before exploring the equity premium produced by the model, we report the performance of the model in accounting for the structural transformation in the data. Figure 2 shows the evolution of the two sectors in time. The share of services in consumption in U.S. data grows from about 40% of GDP in 1950, to almost 70% in 2018. The model accounts well for the initial and final points of the transion, as reported in Figure 2. However, it displays a substantially more concave pattern of the share of services with respect to the data.

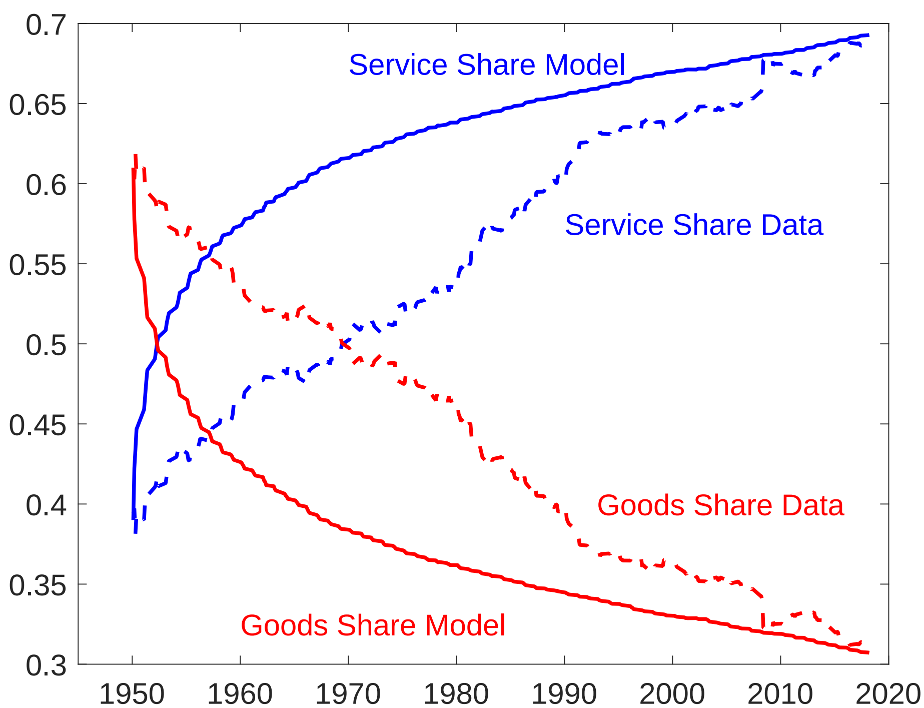

Figure 2: Consumption Shares in Data and Model.

Figure 3 graphically represents our findings for the equity premium in the model. In 1950, the model can account for between 1/3 and 1/4 of the equity premium we measure in the data and report in Figure 1 (4.5% versus 16%). Over time, as the share of services in the economy grows, the equity premium in the economy evolves. In particular, the equity premium exhibits a marked decline at the early stages of the transition from the goods sector

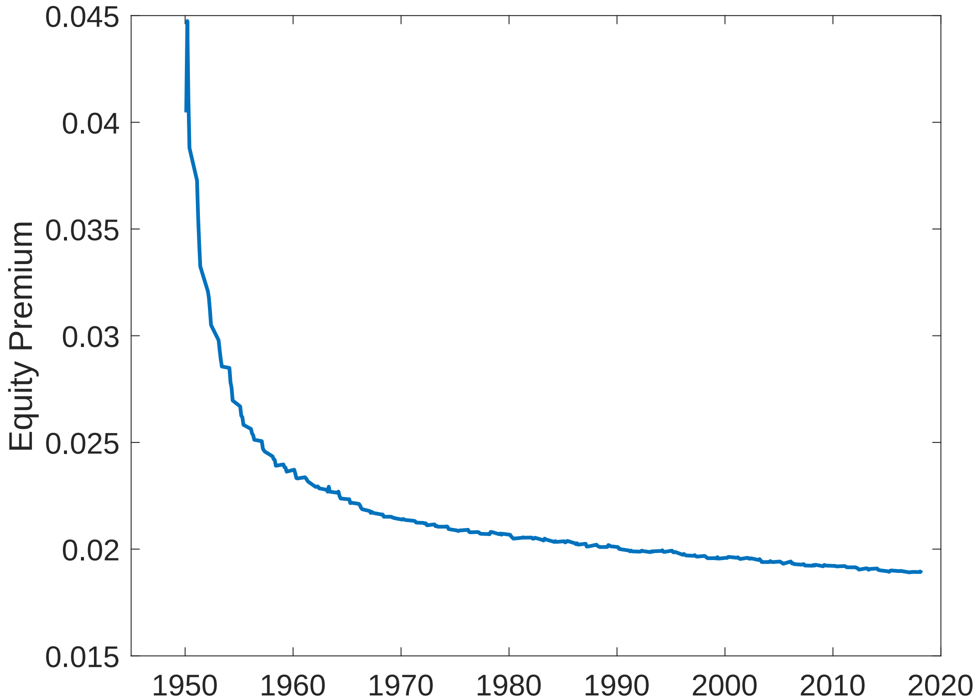

Figure 3: The Equity Premium in the Model

to the services sector and a stabilisation afterwards. This decline of roughly 50% is broadly consistent with the U.S. evidence reported in section 2.

5 Conclusions

[TBA]

References

Ait-Sahalia, Y., Parker, J. A., and Yogo, M. (2004). Luxury goods and the equity premium. *The Journal of Finance*, 59(6):2959–3004.

Boppart, T. (2014). Structural change and the kaldor facts in a growth model with relative price effects and non-gorman preferences. *Econometrica*, 82(6):2167–2196.

Campbell, J. Y. (2008). Estimating the equity premium. *Canadian Journal of Economics/Revue canadienne d'économique*, 41(1):1–21.

Carvalho, V. and Gabaix, X. (2013). The great diversification and its undoing. *The American Economic Review*, 103(5):1697–1727.

Cochrane, J. H., Longstaff, F. A., and Santa-Clara, P. (2008). Two trees. *The Review of Financial Studies*, 21(1):347–385.

Garleanu, N., Panageas, S., and Yu, J. (2012). Technological growth and asset pricing. *The Journal of Finance*, 67(4):1265–1292.

Herrendorf, B., Rogerson, R., and Valentinyi, A. (2014). Growth and structural transformation. *Handbook of Economic Growth*.

Imbs, J. and Wacziarg, R. (2003). Stages of diversification. *American Economic Review*, 93(1):63–86.

Jordà, Ò., Knoll, K., Kuvshinov, D., Schularick, M., and Taylor, A. M. (2019). The rate of return on everything, 1870–2015. *The Quarterly Journal of Economics*, 134(3):1225–1298.

Kogan, L. and Papanikolaou, D. (2012). Economic activity of firms and asset prices. *Annu. Rev. Financ. Econ.*, 4(1):361–384.

Koren, M. and Tenreyro, S. (2007). Volatility and development. *The Quarterly Journal of Economics*, 122(1):243–287.

Koren, M. and Tenreyro, S. (2013). Technological diversification. *American Economic Review*, 103(1):378–414.

Leon-Ledesma, M. and Moro, A. (2020). The rise of services and balanced growth in theory and data. *American Economic Journal: Macroeconomics*.

Lettau, M., Ludvigson, S. C., and Wachter, J. A. (2008). The declining equity premium: What role does macroeconomic risk play? *The Review of Financial Studies*, 21(4):1653–1687.

Martin, I. (2013). The lucas orchard. *Econometrica*, 81(1):55–111.

Moro, A. (2012). The structural transformation between manufacturing and services and the decline in the us gdp volatility. *Review of Economic Dynamics*, 15(3):402 – 415.

Moro, A. (2015). Structural change, growth, and volatility. *American Economic Journal: Macroeconomics*, 7(3):259–94.

Papanikolaou, D. (2011). Investment shocks and asset prices. *Journal of Political Economy*, 119(4):639–685.

Rubini, L. and Moro, A. (2019). Stochastic structural change.