University of Toronto
Department of Economics

Working Paper 516

Towards a Micro-Founded Theory of Aggregate Labor Supply

By Andres Erosa, Luisa Fuster and Gueorgui Kambourov

July 14, 2014

Towards a Micro-Founded Theory of Aggregate Labor Supply*

Andrés Erosa†
Universidad Carlos III de Madrid

Luisa Fuster‡
Universidad Carlos III de Madrid

Gueorgui Kambourov§
University of Toronto

Abstract

We build a heterogeneous life-cycle model which captures a large number of salient features of individual labor supply over the life cycle, by education, both along the intensive and extensive margins. The model provides an aggregation theory of individual labor supply, firmly grounded on individual-level micro evidence, and is used to study the aggregate labor supply responses to changes in the economic environment. We find that the aggregate labor supply elasticity to a transitory wage shock is 1.75, with the extensive margin accounting for 62% of the response. Furthermore, we find that the aggregate labor supply elasticity to a permanent-compensated wage change is 0.44.

JEL Classification: D9, E2, E13, E62, J22.

Keywords: Aggregate labor supply, intensive margin, extensive margin, life cycle, fixed cost of work, non-linear earnings, precautionary savings, preference heterogeneity.

---

*This version: July 4, 2014. We thank the Editor and three anonymous referees for helpful comments, as well as participants in numerous seminars and conferences. This research was supported by a Marie Curie International Reintegration Grant within the 7th European Community Framework Programme and the Social Sciences and Humanities Research Council of Canada. We thank Hui Xiong for excellent research assistance.

†Universidad Carlos III de Madrid, Calle Madrid 126, 28903 Getafe, Madrid, Spain. E-mail: aerosa@eco.uc3m.es.

‡Universidad Carlos III de Madrid, Calle Madrid 126, 28903 Getafe, Madrid, Spain. E-mail: lfuster@eco.uc3m.es.

§University of Toronto, 150 St. George St., Toronto, ON, M5S 3G7 Canada. E-mail: g.kambourov@utoronto.ca.

1 Introduction

How responsive is aggregate labor supply to changes in the economic environment? While the answer to this question is crucial for understanding the effects of government policies and business cycles, it has led to substantial debate in the economics profession. At the core of this debate is the inconsistency between the small labor supply elasticities estimated in micro studies and the large ones used in macro models. The empirical evidence indicates that aggregate labor supply responses are determined by individual responses along both the intensive and the extensive margins.[1] Further, economic theory implies that labor supply responses along the intensive and extensive margins are distinct objects. The intensive margin responses are mainly driven by the intertemporal substitution of labor (the Frisch elasticity of labor supply). The extensive margin responses, on the other hand, are unrelated to the preference parameter typically estimated in micro studies. Chang and Kim (2006) build on the insights from the model of indivisible labor in Hansen (1985) and Rogerson (1988), introduce heterogeneity, and show that the slope of the aggregate labor supply schedule is determined by the distribution of reservation wages rather than by the willingness to substitute leisure intertemporally, establishing that when the extensive margin is operative heterogeneity and aggregation play a crucial role in determining aggregate labor supply responses. Therefore, the starting point for developing a theory of aggregate labor supply behavior, which can be used with some confidence in analyzing the effects of various government policies and macroeconomic shocks, involves building a rich quantitative theory of heterogeneity that is consistent with individuals' labor supply behavior along many dimensions. This is the goal of this paper.

The literature modeling labor supply decisions has long recognized that many factors affect the mapping from individual preference parameters to labor supply responses at the macro level, as discussed in recent surveys by Keane (2011) and Keane and Rogerson (2011, 2012). Motivated by the discussion in this literature, we develop a quantitative life-cycle theory to study aggregate labor supply responses in a framework that incorporates the following features: (i) wage risk, (ii) employment risk and labor market frictions, (iii) heterogeneity in tastes and skills, (iv) tied wage-hours offers, (v) fixed costs of work, (vi) credit constraints and incomplete markets, and (vii) government programs (tax and transfer programs). Our contribution is to integrate all of these features into a unified framework of labor supply,

[1] See Cooley (1995) for evidence on the adjustment in labor supply along both margins over the business cycle and Blundell et al. (2011) for recent evidence on the importance of both margins over time in the US, the UK, and France.

argue that they are all important for understanding individual-level data on labor supply, and analyze their importance for aggregate labor supply responses to changes in the economic environment.

We discipline our theory with a large number of facts on labor supply at the individual level documented on data from the Survey of Income and Program Participation (SIPP) and the Panel Study of Income Dynamics (PSID). Most importantly, the SIPP data reveals a quite operative extensive margin at high frequencies, such as 4-month periods referred to as *quadrimesters*: the probability of entering non-employment is high, most non-employment spells are short-lived, and this pattern is observed at all stages of the life cycle. Thus, individuals frequently use the extensive margin to adjust their labor supply within a year, even at old ages. Capturing this pattern is central to the analysis in this paper, and we find that it is essential for understanding labor supply responses. In addition, the PSID data reveals another insightful pattern: there are large differences in lifetime labor supply across individuals, even conditional on education and permanent income, leading to our finding that preference heterogeneity is an important factor for understanding the variation in labor supply across individuals.

Motivated by the evidence in Aaronson and French (2004), we assume that earnings are a convex (non-linear) function of hours of work. The convexity in the returns to work implies that individuals do not have incentives to work short hours, making the extensive margin of labor supply decisions more important. We introduce employment risk into the framework because the SIPP data indicates that a non-trivial fraction of the entries into non-employment represents involuntary separations. The calibration implies that individuals face substantial fixed costs of work, a job finding rate of less than one when non-employed, and that there is significant preference heterogeneity in the taste for leisure. All of these features are needed for the model to be consistent with the following moments: (i) the age-profile of employment rates, (ii) the probability of transitioning from employment to non-employment at various stages in the life cycle, (iii) the duration distribution of non-employment spells, and (iv) the lifetime inequality in labor supply. Thus, our calibration strategy ensures that the theory is consistent with the incidence and volatility of non-employment spells and the lifetime dispersion in labor supply − an important step in the development of a quantitative theory of labor supply responses.

We emphasize that the life-cycle patterns implied by the theory are the result of the mechanisms embedded in the model. The accumulation of assets over the life cycle is important for the theory to match various non-targeted life-cycle facts. First, while the calibration targeted statistics on the duration distribution of *all* non-employment spells, the model captures quite well how the duration distribution of non-employment spells varies over the life cycle. When individuals are young, about 60% of the non-employment spells last one period (quadrimester). As individuals age and accumulate assets, this fraction decreases to about 40% for the 55-61 age group. On the other hand, the fraction of non-employment spells with a length of at least four periods grows over the life cycle reaching a value close to 40% for the age group 55-61. Overall, we show that our baseline economy captures very well the dynamics of labor supply at the extensive margin for various age and education groups. Second, as in the data, the model economy predicts that the cross-sectional variance in hours of work increases with age. Third, the baseline economy matches quite closely how lifetime labor supply (the coefficient of variation over the sum of hours worked during a 10 year period) varies by age groups in the PSID. As individuals age and build precautionary savings, the heterogeneity in asset holdings translates into more persistent differences in labor supply across individuals. Fourth, consistent with the evidence, incomplete markets imply that the variance of consumption grows about 12 log points over the life cycle in the baseline economy.

We find that the aggregate labor supply elasticity to a one period unanticipated small wage change in the baseline economy is 1.75, with an extensive margin elasticity of 1.08 (about 62% of the aggregate labor supply response). Moreover, the aggregate response is much higher for high school individuals than college educated individuals (1.96 versus 1.35). Through a series of experiments aimed at assessing the importance of various modeling assumptions, we conclude that even though the intensive-margin elasticity of aggregate labor supply is quantitatively important it is almost invariant to changes in the modeling assumptions. The effects of the modeling assumptions on labor supply responses are almost entirely driven by their impact on the extensive margin. When we shut down the non-linear earnings feature in the baseline economy and simulate the employment responses to a small temporary wage change, we find that the fixed costs of work in the baseline economy are sufficiently small that, on their own, they do not affect employment decisions along the extensive margin. However, fixed costs of work matter importantly in the presence of non-linear earnings, accounting for half of the employment response in our baseline economy: When we shut down fixed costs of work in the baseline economy, the extensive margin elasticity to a temporary wage change drops from 1.08 to 0.57. This result underscores that there is a *complementarity* between fixed costs of work and non-linear earnings that *enhances* the aggregate labor supply response to a temporary wage change.

The extensive margin is also crucial for understanding how the labor supply response to a temporary wage change varies across age and education groups. The age profile of the labor supply response in the baseline economy is U-shaped. Such a life-cycle pattern is reminiscent of the fact that over the business cycle hours of work and employment fluctuate much more for young and old individuals than for the middle-aged, as discussed in Gomme et al. (2004) and Jaimovich and Siu (2009). Modeling credit (or initial assets), preference heterogeneity, and the extensive margin are important for generating the decline in the labor supply response early in the life cycle. If any of these features are shut down in the baseline economy, the elasticity of labor supply does not decrease at young ages. Modeling productivity risk and the extensive margin (non-linear earnings) are important for accounting for the increase in the labor supply elasticity to a temporary wage change late in the life cycle. Preference heterogeneity increases aggregate labor supply responses, and this effect is the strongest for young individuals with high school education. Finally, both labor market risk and labor productivity risk have large effects on employment decisions of old individuals (aged 55-61) but these effects have opposite sign. Labor market frictions make old individuals less willing to take periods off work, so that they respond less to temporary wage changes. On the other hand, productivity risk encourages individuals to build precautionary savings over the life cycle increasing their labor supply response to temporary wage changes at old age.

Regarding permanent (compensated) wage changes, we find that the Hicks elasticity of aggregate labor supply in the baseline economy is 0.44.[2] Interestingly, labor supply responses along the intensive and extensive margins have different signs. While the intensive margin elasticity is 0.55, the extensive margin elasticity is -0.11. Hours along the intensive margin decrease because of the substitution effect associated with the lower wage rate. On the other hand, the decrease in the (compensated) wage discourages savings over the life cycle which, in turn, leads to an important increase of the employment rate at old age. This effect accounts for the aggregate negative elasticity of labor supply along the extensive margin. We also find that in the economy with linear earnings the aggregate Hicks elasticity is 0.56, which is higher than the 0.44 estimated for the baseline economy. The Hicks elasticity is now higher than in the baseline economy because the extensive margin elasticity is zero instead of negative. Shutting down preference heterogeneity also increases the aggregate Hicks elasticity from 0.44 in the baseline economy to 0.49. Hence, unlike the findings for temporary wage changes,

[2]The mean value of the Hicks elasticity of aggregate hours across the micro studies reviewed in Chetty et al. (2011) is 0.76, with substantial variation in the estimated elasticities across studies. Since there are wide confidence intervals associated with each of the point estimates as well as methodological disputes about the validity of some of the studies, Chetty et al. (2011) argue that the estimates should be treated as rough values meant to gauge the order of magnitudes. In this sense, the Hicks elasticity implied by our model appears reasonable.

preference heterogeneity does not matter much for the Hicks elasticity of labor supply.

We also re-calibrate a version of the model economy with linear earnings and compare it to the baseline economy, both in terms of their fit of the micro facts as well as in terms of their implications for aggregate labor supply responses. We find that the economy with linear earnings implies a worse fit of the calibration targets. The most striking difference is that the economy with linear earnings requires fixed costs of work about four times higher than in the baseline economy: for instance, focusing on high school individuals, the economy with linear earnings requires a fixed cost of work equivalent to a consumption loss of 36%, instead of the 8.5% value in the baseline economy. In the paper we discuss evidence on work-related expenditures from Aguiar and Hurst (2013) suggesting that the fixed costs of work in the baseline economy are empirically more plausible.[3]

We also compare the response to a one period (quadrimester) wage change of 2% across the two calibrated model economies. In the re-calibrated economy with linear earnings the elasticity of aggregate labor supply along the extensive margin is less than a half the one in the baseline economy (0.44 versus 1.08). We draw two key lessons: First, fixed costs of work may allow for significant (bounded away from zero) extensive margin responses if they are large enough. Second, non-linear earnings *amplify* the effect of fixed costs of work on extensive margin responses to a temporary wage change. The baseline economy (with non-linear earnings) has an extensive margin response to a 2% wage increase that is 2.5 times higher (e.g. the ratio between 1.08 and 0.44 is about 2.5) than the re-calibrated economy with linear earnings even though fixed cost of work are about a factor of four smaller in the former economy.

There is a large literature in macroeconomics analyzing labor supply. To the best of our knowledge, none of the papers in the macro literature has jointly studied all of the features incorporated into our framework. To name a few, Domeij and Flodén (2006) and Pijoan-Mas (2006) study labor supply decisions in an incomplete markets framework with borrowing constraints but they abstract from life-cycle behavior, preference heterogeneity, and the modeling of the extensive margin. Low (2005) does model life cycle and incomplete markets but abstracts from the extensive margin. On the other hand, Chang and Kim (2006) model the extensive margin but abstracts from labor supply decisions along the intensive margin, and they do not model preference heterogeneity, life cycle, and tied wage offers (non-linear earn-

[3]Another important difference across the calibrated model economies is that the baseline economy features much lower preference heterogeneity among non-college individuals. The standard deviation of the fixed effect on taste for leisure in the baseline economy is 0.12 whereas it is 0.18 in the economy with linear earnings. In addition, we find that the economy with linear earnings implies a counterfactually high inequality in the distribution of hours of work and in lifetime hours, and too high variance of consumption.

ings). Relative to Rogerson and Wallenius (2009), our contribution is to build a theory of aggregation, disciplined with micro data, with several dimensions of heterogeneity (skills and tastes). Moreover, by modeling non-linear earnings *together* with fixed costs of work, we can focus on how these two features interact in determining labor supply responses along the extensive margin. Perhaps the closest paper to us is French (2005), who incorporates most of the features in our paper but models the extensive margin at an annual level because he is mostly interested in studying retirement decisions late in the life cycle. Finally, our framework abstracts from human capital accumulation, a feature that Imai and Keane (2004) and Keane (2011) have emphasized. Instead, our analysis focuses on different, but complementary, mechanisms.

The paper is structured as follows. Section 2 presents empirical facts on labor supply that guide the development of our theory. Section 3 develops the theory and provides some analytical insights about how non-linear earnings, coupled with fixed costs of work, affect labor supply decisions along the intensive and extensive margins. The calibration of the model economy, the calibration results, and performance of the baseline economy along non-targeted dimensions are discussed in Section 4. Section 5 studies aggregate labor supply responses to temporary and permanent (compensated) wage changes and discusses how various model features impact on aggregate labor supply elasticities. We show that labor supply responses are consistent with evidence from a tax reform in Iceland in 1987, as discussed in Bianchi et al. (2001).[4] Section 6 analyzes the performance of two alternative recalibrated economies: an economy with linear earnings and an economy with a lower intertemporal elasticity of substitution of leisure. Section 7 concludes.

2 Empirical facts

In this section we describe the facts on labor supply over the life cycle, lifetime labor supply, and evidence on non-linear earnings that guide the development of our theory.

2.1 Labor supply over the life cycle

We begin our empirical analysis by describing facts on labor supply at the extensive (whether to work or not) and intensive (how much to work) margins. The distinction between the extensive and intensive margins of labor supply depends on the period of time. In our analysis we will use data from the Survey of Income and Program Participation (SIPP), and we will

[4]In a recent survey, Chetty et al. (2011) claim that state-of-the-art macro models are grossly at odds with evidence on the effects of the tax reform in Iceland in 1987.

define the time period to be 4 months (a quadrimester) implying that variation in labor supply within a quadrimester period is interpreted as changes along the intensive margin while variation in the number of quadrimesters worked would be interpreted as changes in labor supply along the extensive margin.

The SIPP interviews individuals three times a year (once every four months) and provides detailed monthly information on labor market history and income and welfare program participation. We use this information in oder to compute labor market statistics of interest at a high frequency, such as a quadrimester. We use the 1990 SIPP Panel which runs from October 1989 until August 1992. That implies that it spans 8 quadrimesters (32 months). We restrict the analysis only to men, who are not self-employed, between the ages of 25 and 61. Individuals are classified as non-college if they have either elementary or high-school education and as college if they report having completed college education.

Definition of employment and non-employment. In order to classify an individual as either employed or non-employed in a particular quadrimester, we utilize the available information in the SIPP on hours worked at the monthly level.

First, we identify whether an individual in a particular month is (1) working (i.e., reporting positive hours worked), (2) not working (i.e., reporting zero hours worked), or (3) a non-respondent, not in the sample. Only months with information about working (i.e. either (1) or (2)) can be used in the analysis. Individuals who work as self-employed in a given month fall into category (3) in that particular month. Second, if an individual works in one or more months in a given quadrimester, regardless of whether the remaining months fall into categories (2) or (3), then he is classified as employed in that quadrimester. If an individual falls into category (2) − i.e., reporting zero hours worked − in each of the four months in a quadrimester, then he is classified as non-employed. All other individuals are dropped from the sample for that quadrimester. Finally, a non-employment spell starts whenever an individual is working in one quadrimester and not in the next one. The non-employment spell ends when the individual records a quadrimester of employment.

Employment rate. Figure 1 shows the quadrimesterly employment rates for college and non-college individuals over the life cycle. In particular, we group individuals in 4 age groups: 25-34, 35-44, 45-54, and 55-61. The employment patterns are not surprising: college individuals have a higher employment rate, at all ages, than non-college individuals, and there is a pronounced decline in the employment rates late in the life cycle.[5]

[5]Note that measures of employment that use information on hours worked in the week preceding the interview will produce slightly lower employment rates than our definition of employment.

Figure 1: Employment rate, 1990 SIPP.

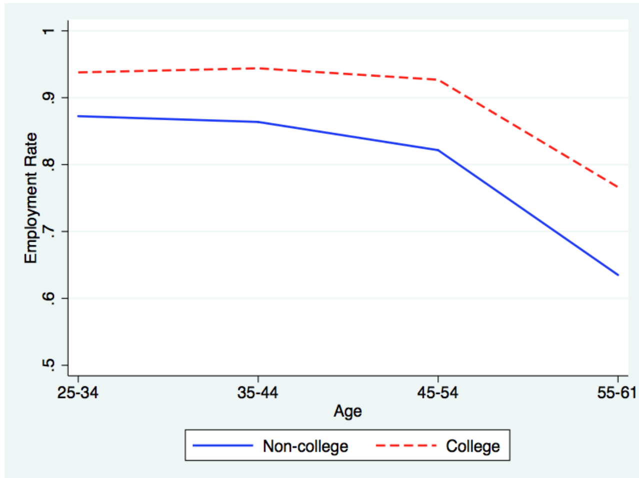

Figure 2: Mean hours worked, 1990 SIPP.

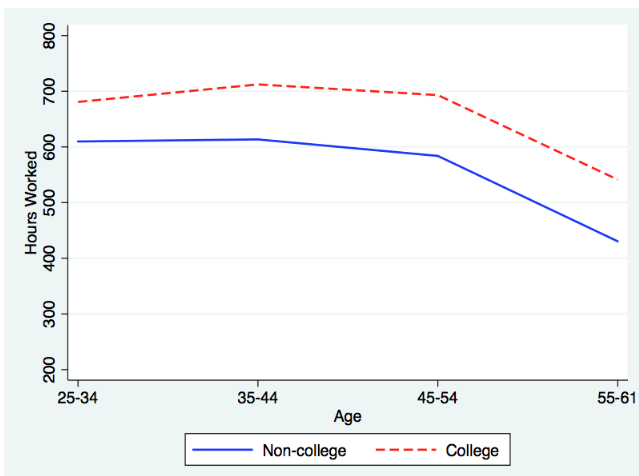

Mean hours worked. Figure 2 reports mean hours worked at the quadrimesterly level, by age and education groups. Average hours worked are higher for college individuals, and mean hours worked decrease late in the life cycle for both education groups. This is due to both lower employment rates and lower hours per worker, although the extensive margin is quantitatively much more important.

Durations of non-employment spells. Figure 3 displays the histograms of non-employment spells for non-college and college individuals. These histograms provide us with important information regarding the nature of labor supply at the extensive margin. Approximately half of all non-employment spells last for only one quadrimester, and only around 25% of the spells

Figure 3: Non-employment spells, 1990 SIPP.

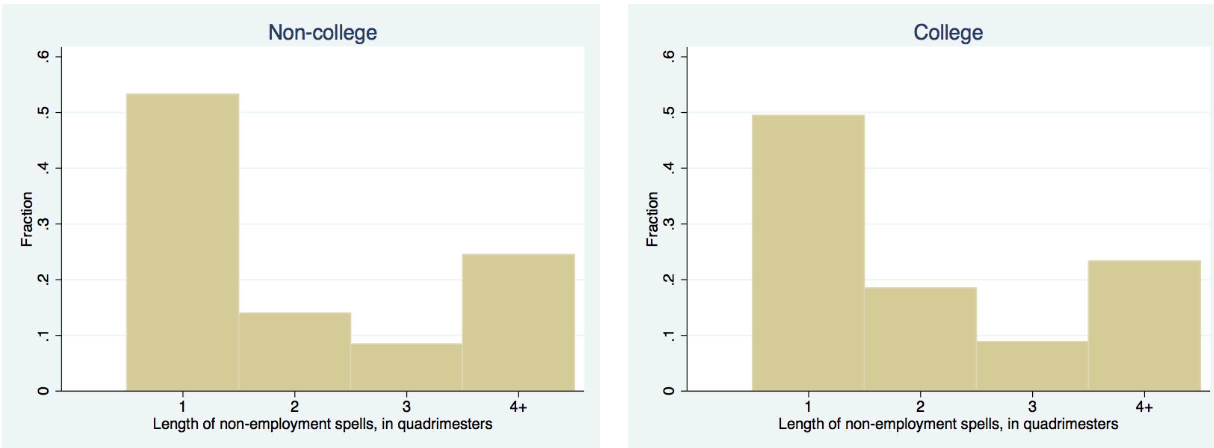

last for more than 4 quadrimesters (i.e., more than 16 months). In general, college individuals tend to spend more periods in non-employment, although as we will see below they are much less likely to enter non-employment spells.

Figure 4: Non-employment spells, 1990 SIPP: by age and education.

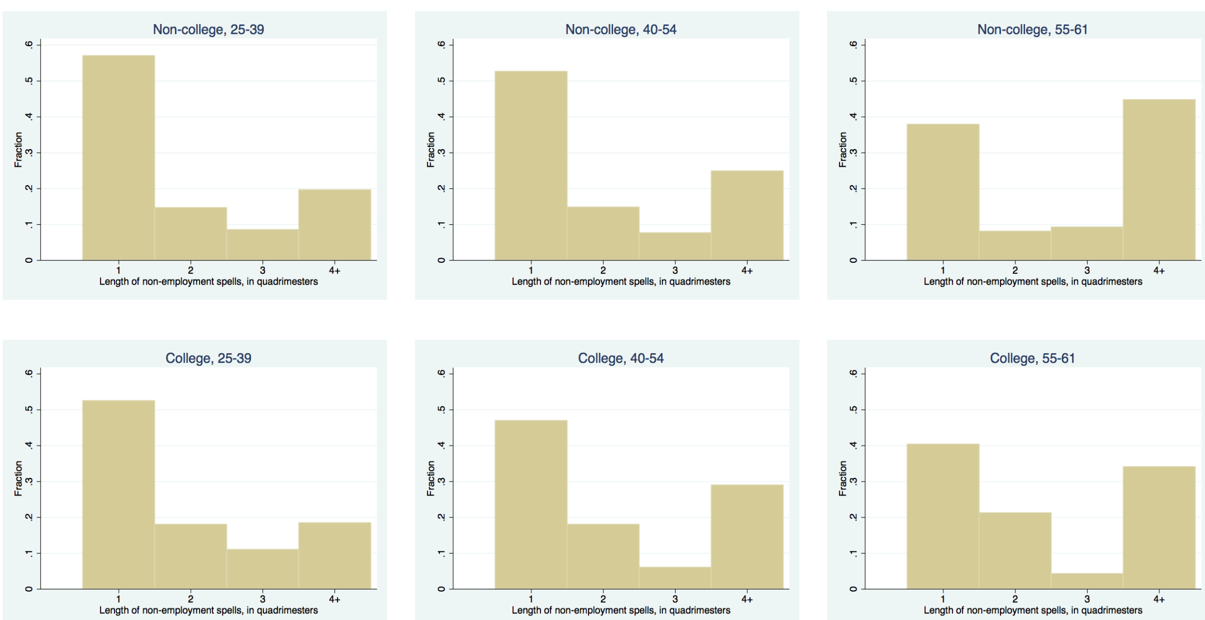

Figure 4 reveals additional insightful information by reporting the histograms of non-employment spells also by age. We observe a pronounced life-cycle pattern, with older individuals spending more periods into non-employment. Nevertheless, even for the old age group of 55-61 year olds, around 50% (60%) of the non-employment spells for non-college (college) individuals last for one or two periods, indicating an active extensive margin of labor supply even at an old age.

Entry into non-employment spells. Figure 5 reports the fraction of those entering a non-employment spell, by age and education groups. Non-college individuals are much more likely to enter a non-employment spell, and that fraction increases late in the life cycle. Recall that individuals are classified as non-employed if they do not work at all throughout the whole 4-month period. Therefore, those who use the extensive margin not to work spend a non-trivial amount of time non-employed. Nevertheless, we find very large entry rates into non-employment, at all ages of the life cycle, indicating that individuals consistently use the extensive margin in order to adjust their labor supply: throughout most of the life cycle around 4% (2%) of the employed non-college (college) individuals will find themselves non-employed in the following quadrimester. These numbers increase to 6% and 4% late in the life cycle for the non-college and the college, respectively.

Figure 5: Entry rate into non-employment, 1990 SIPP.

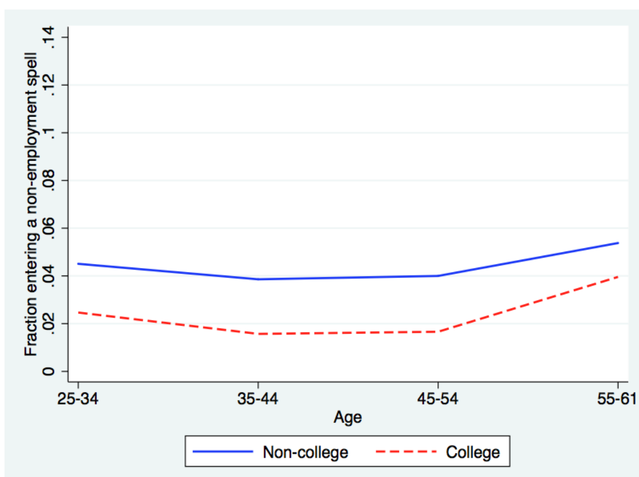

Notes: The entry rate into non-employment is measured as the fraction of those employed in quadrimester t that are non-employed in quadrimester $t+1$.

We use SIPP data on transitions from employment into unemployment in order to differentiate labor market transitions out of employment that might be due to involuntary separations. In order to analyze entry into unemployment, we focus on those individuals who report to have been employed in the previous quadrimester, but are classified as non-employed in the current quadrimester according to the criteria specified above. We then use three variables from the 1990 SIPP to identify the subset of individuals from the group of non-employed who are unemployed: (i) number of weeks looking for work, (ii) employment status recode for each month, and (iii) amount of unemployment compensation and benefits. We consider an individual to be unemployed if he spent any positive time looking for a job or received unemployment benefits and compensation. The resulting quadrimesterly entry rates into unemployment are, as expected, higher for non-college than college individuals. Moreover, for both education groups they are fairly stable over the life cycle: 2.4% for non-college and and 1.16% for college. Hence, our findings suggest that about half of the non-employment spells are voluntary. Nonetheless, we emphasize that these numbers represent a lower bound on the actual amount of voluntary separations in the data for two reasons. First, we consider an individual to be non-employed only if the non-employment spell lasts for at least 4 months, thereby not capturing short non-employment spells. Second, our procedure most likely underestimates voluntary separations because an individuals who voluntarily separates from a job in a given quadrimester and searches for a job (at some point) during the next quadrimester is counted as unemployed. Hence, if anything, our procedure minimizes the importance of extensive margin responses.

2.2 Lifetime labor supply

The dispersion in lifetime labor supply is another useful statistic which is closely related to the persistence in an individual's labor supply over time. We use the Michigan Panel Study of Income Dynamics (PSID) for the period 1968-1997 in order to compute all annual statistics. The sample is restricted to men between the ages of 25 and 61. We consider an individual to be high school if he has at most 12 years of education while those with 14 years of education or more are considered to be college graduates. We keep in the sample those born between 1922 and 1962.

Due to the nature of the PSID dataset, we do not observe individuals throughout all their life − some of them have already been in the labor market for some time when the survey starts in 1968 while those who enter the labor market in 1968 at the age of 25 are only in their 50s in 1997. Nevertheless, we can learn a lot even if we follow individuals for shorter periods. We choose to follow individuals for periods of 10 years at different stages in their life cycle: ages 25-34, 35-44, 45-54, and 55-61. We drop all individuals who have a missing observation during the relevant ten years and sum the hours worked for each individual during the whole ten years. Then we compute the coefficient of variation in this cumulative measure of hours worked and refer to it as a measure of the dispersion in lifetime hours. In addition, we also compute, on exactly the same sample, the coefficient of variation of cross-sectional annual

Figure 6: Coefficient of variation in hours and lifetime hours, PSID.

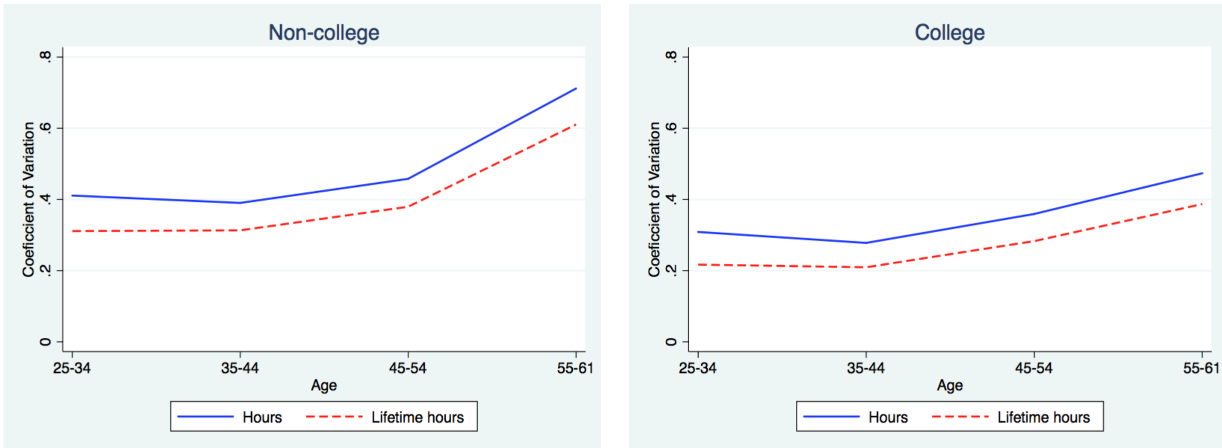

Notes: For each ten year period (e.g., between the ages of 35 and 44), let h_{ij} denote the annual hours worked for individual i who is j years old and let $\tilde{h_i} = (\sum_{j=35}^{44} h_{ij})$ denote the individual i's lifetime hours during this period. The coefficient of variation in lifetime hours across all i individuals is $CV(\tilde{h_i})$. The coefficient of variation in hours across all individuals of age j is $CV^j(h_{ij})$. The figure reports the average $CV^j(h_{ij})$ over the ages 35-44, $E_j[CV^j(h_{ij})]$.

hours and refer to it as a measure of the dispersion in hours.[6]

Considering two extreme examples is useful for illustrating how one can interpret these two measures of the dispersion in hours. Consider a particular group, e.g. the group between the ages of 35 and 44, and suppose that all individuals work the same number of hours throughout the whole period as at the beginning at the age of 35. In that case, the coefficient of variation of the cumulative hours worked throughout the whole period would be the same as the coefficient of variation (cross-sectionally) at the age of 35 (or any other age in the period). Alternative, suppose that individual hours fluctuate a lot over the period and those who work a lot in one year work very little the year after that. In that case, workers would end up working quite similar cumulative hours over the period, and the coefficient of variation of the cumulative hours worked throughout the whole period would be quite small and substantially lower than the coefficient of variation (cross-sectionally) at the age of 35 (or any other age in the period).

Figure 6 reports the measured dispersion in cross-sectional and lifetime hours for the four age groups defined above. Non-college individuals display higher dispersions in both cross-

[6]In particular, for each ten year period (e.g., between the ages of 35 and 44), let h_{ij} denote the annual hours worked for individual i who is j years old and let $\tilde{h_i} = (\sum_{j=35}^{44} h_{ij})$ denote the individual i's lifetime hours during this period. The dispersion in lifetime hours is measured by $CV(\tilde{h_i})$, the coefficient of variation in lifetime hours across all i individuals. The cross-sectional dispersion in hours at age j is measured by $CV^j(h_{ij})$, the coefficient of variation in hours across all individuals of age j. The figure reports the average $CV^j(h_{ij})$ over the ages 35-44, $E_j[CV^j(h_{ij})]$.

sectional and lifetime hours. Further, the analysis provides us with two important findings. First, the dispersion in cumulative hours is quite substantial, indicating that individuals tend to be quite persistent in their labor supply behavior. Second, the dispersion of cumulative hours is smaller than the cross-sectional dispersion at any age in the 25-61 interval. This implies that workers do sometimes change their hours worked.

Lifetime labor supply for high- and low-productivity individuals. Using the fact that the PSID is a long panel we show that there is no association between lifetime labor supply and average labor productivity (wages) over the life cycle across individuals of the same education group. To this end, for each cohort and education group we divide individuals into high and low productivity types. First, we compute each individual's mean wage over the age of 30 to 45 and classify them into high and low types depending on whether their mean wages are above or below the median wage in their cohort-education category. We then compute mean hours worked for high and low types and find that there are virtually no differences in labor supply. Focusing on the age group 30-45 and non-college individuals, the average hours worked across all cohorts is 2143 for type 1 individuals and 2166 for the type 2. For individuals with college education, average hours are 2269 and 2271 for type 1 and type 2, respectively.

2.3 Non-linear earnings

In cross-sectional data, there is a positive correlation between hourly wages and hours worked − i.e., those that work more hours also receive higher hourly wages. This result could be due to selection (those who work long hours may be more productive) or could be due to the fact that individuals choose to work less hours during periods of relatively low hourly wages. However, various studies have also documented the existence of an additional mechanism − individuals that decide (for reasons other than changes in the hourly wage) to work fewer hours do get offered lower hourly wages.[7] For example, Gustman and Steinmeier (1985) document that individuals that enter partial retirement experience − depending on whether they remain in their main job or not − a 15-23% drop in their hourly wages.[8] More recently, Aaronson and French (2004) estimate the quantitative importance of this mechanism by exploiting the sharp decline in hours worked for working men at the ages of 62 and 65 in US data. They treat this decline in hours worked for those who work (relative to a quadratic age polynomial

[7] See Gustman and Steinmeier (1985) and Aaronson and French (2004) for references to the literature that has studied this issue.

[8] Gustman and Steinmeier (1986) argue that a structural model of retirement with such a feature matches very well the labor supply patterns after the age of 60.

in hours worked over the life cycle) as exogenous variation in hours worked, and thus a useful instrument in the empirical analysis, caused by the US social security rules that discourage individuals at those ages from working long hours. They find that men who cut their work-week from 40 to 20 hours experience a 20-25% decline in their hourly wage. In order to control for the fact that workers that reduce their working hours might be also changing the type of work they do on the job, either by switching employers or by switching their occupation within the same employer, they further restrict the analysis to workers that remain in the same job and/or same occupation. Finally, since the various datasets used in the analysis are all longitudinal and follow individuals over time, the estimation controls for individual fixed effects and thus for unobserved quality and productivity differences between those that work long and short number of hours.[9]

3 Model

We develop a life-cycle theory of labor supply of individuals along the intensive and extensive margins. The model abstracts from the labor supply decisions of women and analyzes only males. We consider a small open economy facing a fixed interest rate.

3.1 Population, preferences, and endowments

The economy is populated by overlapping generations of individuals who start their lives at age 25, face uncertain lifetimes, and live, at most, J periods. They differ in terms of their education (college versus non-college), labor productivity, and taste for leisure. The college decision is exogenous, and the education type of an individual determines the stochastic processes driving the mortality, taste, and labor productivity shocks.

Preferences. The date-t utility function takes the form

$$u_t = u(c_t, l_t) = (1-\varphi)\ln c_t + \varphi \frac{l_t^{1-\sigma}}{1-\sigma} - I_{(l_t<1)}F, \tag{1}$$

where c_t is consumption and l_t denotes leisure. Individuals are heterogeneous in their taste for leisure φ, which evolves stochastically over time. Individuals who work ($l_t < 1$) face a fixed disutility cost of work (F). The utility function is consistent with balanced growth − this assumption allows the theory to be consistent with the fact that there are large permanent

[9]Moffitt (1984), using taste shifters, such as children, as instruments in a wage on hours regression, and Keane and Wolpin (2001), using a structural model, also find evidence of tied wage-hours offers.

differences in labor productivities across individuals (heterogeneity in fixed effects) but not in their lifetime labor supply, as discussed in Section 2.2.

Individuals maximize lifetime expected utility

$$E \sum_{t=1}^{J} \beta^t \pi_t \, u(c_t, l_t), \tag{2}$$

where π_t denotes the probability that the individual survives to age t and E denotes the expectation operator. Each period, as described below, individuals face mortality shocks, labor market risk (job separation and job finding risk), shocks to the taste for leisure (φ), and labor productivity risk (z).

Heterogeneity in preferences. For each education group, it is assumed that taste shocks follow the stochastic process

$$\varphi_{it} = \varphi_i \times \varphi_t,$$

where $\varphi_i \sim N(\mu_\varphi, \sigma_\varphi^2)$ is an individual fixed effect determined at birth. The term φ_t represents a stochastic deviation from the mean value φ_i, which follows a first-order autoregressive process.

$$\begin{aligned} \varphi_t &= \rho_\varphi \varphi_{t-1} + \eta_{\varphi_t}, \;\; \eta_{\varphi_t} \sim N\left(1 - \rho_\varphi, \sigma_{\eta_\varphi}^2\right), \\ \varphi_0 &\sim N\left(1, \frac{\sigma_{\eta_\varphi}^2}{1 - \rho_\varphi^2}\right), \end{aligned} \tag{3}$$

where ρ_φ denotes the persistence of the shock on preferences for leisure and η_{φ_t} is the innovation at age t. When individuals enter the model economy, the initial seed for the autoregressive process (φ_0) is drawn from the invariant distribution of preference shocks. This assumption implies that the distribution of shocks does not vary over the life cycle, so that the life-cycle patterns of labor supply implied by the theory are not due to variation in preferences over the life cycle. While the parameters $(\mu_\varphi, \sigma_\varphi^2, \rho_\varphi, \sigma_{\eta_\varphi}^2)$ vary across education types, this is omitted to simplify the notation.

Labor services. An individual's time endowment in each period is one. The amount of time that can be allocated to work is $h_j = 1 - l_j$. It is assumed that the mapping from hours of work to labor services is non-linear, as in Hornstein and Prescott (1993), French (2005), and Rogerson and Wallenius (2009). The idea is that h hours of work map into h^θ units of labor services, with $\theta \geq 1$. The case $\theta = 1$ corresponds to the standard model that assumes a linear mapping from hours to labor services. The case $\theta > 1$ gives rise to a model economy in which earnings are a non-linear function of hours of work. Earnings of an individual with labor productivity z and working hours h are

$$w(z,h) = zh^{\theta}. \tag{4}$$

The calibration of the model will rely on estimates by Aaronson and French (2004) to pin down the value of θ. In Appendix A, we use the production technology proposed by Hornstein and Prescott (1993) to show that it gives rise to a competitive equilibrium in which earnings follow the functional form specified in (4).

Heterogeneity in labor productivity. For each education group, labor productivity z_t is assumed to change stochastically over the life-cycle according to:

$$\ln(z_{it}) = \mathbf{x_t}\kappa + \alpha_i + u_t, \tag{5}$$

where z_{it} denotes labor productivity of individual i at age t, $\mathbf{x_t}$ is a quartic polynomial in age, κ is a vector of coefficients, $\alpha_i \sim N\left(-\frac{\sigma_{\alpha}^2}{2}, \sigma_{\alpha}^2\right)$ is a fixed effect determined at birth, and u_t is a persistent productivity shock. Each period, it is assumed that with probability 2/3 the productivity shock remains constant $u_t = u_{t-1}$ and with probability 1/3 it follows a first-order autoregression:

$$u_t = \rho_u u_{t-1} + \eta_{u_t}, \quad \eta_{u_t} \sim N\left(-\frac{\sigma_{\eta_u}^2}{2}, \sigma_{\eta_u}^2\right), \quad u_0 = 0. \tag{6}$$

While the parameters $\Omega = (\kappa, \sigma_{\alpha}^2, \rho_u, \sigma_{\eta_u}^2)$ vary across education types, this is omitted to simplify the notation. At age 62, labor productivity becomes zero so that all individuals are retired by this age.

Labor market frictions. Individuals start each period in three possible labor market states: employed (e), unemployed (u), non-employed (n). The crucial distinction between the last two states is that unemployment insurance benefits b_u are assumed to be paid only to the unemployed. Employed individuals move into the unemployment state when they receive an exogenous job separation shock, δ. We assume that both unemployed and non-employed individuals face a job finding rate of p. Both δ and p vary across the two education types. The transitions in labor market status across periods depend on the labor supply decisions of individuals (whether to work or not) and on labor market shocks (exogenous job separation and job finding rates), as described in Table 1. Individuals who start the period as employed and choose to work ($h > 0$) face at the end of the period a job separation shock: With probability δ they are exogenously separated and start the next period in the unemployment state, or with probability $(1-\delta)$ they do not suffer a job separation and start the next period as employed. Individuals that start the period as employed and choose not to work ($h = 0$) find a job at the end of the period with probability p so that the next period they start as either employed (with probability p) or non-employed (with probability $1-p$). Since unemployed individuals find a job with probability p, they transit to the employment state with such probability. Unemployed individuals that do not find a job (an event with probability $1-p$), transit to the non-employment state in the next period. Since unemployed individuals transit into either employment (with probability p) or non-employment (with probability $1-p$), individuals can only collect unemployment insurance during the period they were exogenously separated from their job. The parameters δ and p do not vary with age over the life cycle (but do vary across education groups).

Table 1: Labor market transitions.

		s'	
s and hours	e	u	n
e, $h > 0$	$1-\delta$	δ	0
e, $h = 0$	p	0	$1-p$
u	p	0	$1-p$
n	p	0	$1-p$

3.2 Government

The government taxes consumption, capital income, and labor income. It is assumed that consumption and capital income are taxed at flat rates (τ_c, τ_k). Earnings y are taxed according to a progressive tax schedule $T(y)$. The tax revenue is used to finance government expenditures, which are assumed not to provide utility to individuals or to enter in an additively separable fashion in the utility function. The government administers an unemployment insurance scheme that provides unemployment benefits for one period to those individuals who were laid off in the current period. Unemployment benefits are set to a fixed proportion of individual's potential earnings, computed as the earnings the individual would have made had he worked full time − 40% of available time − during the quadrimester.

The government also administers a pay-as-you-go social security system. We model a stylized representation of the US social security system. Pension benefits in the United States are a function of the Average Indexed Monthly Earnings (AIME) over the 35 highest earnings years. Modeling in detail how pension benefits depend on the history of earnings of individuals requires the modeling of a state variable that summarizes average past earnings − an approach which substantially complicates the computational task in a model which is rich in many other dimensions. To simplify the computational procedure, we follow a large literature (see for instance, Low et al. (2010) and Kaplan (2012)) in assuming that pension benefits are a function of individual's fixed characteristics: education and the fixed effect in labor market productivity. We assume that the government collects a payroll tax (τ_{ss}) to finance social security outlays.

3.3 Credit markets

Individuals can insure mortality risk in fair annuity markets.[10] Denoting by R the gross interest rate net of capital income taxes τ_k, the gross interest rate faced by an individual j years old with education e is given by

$$R_j^e = 1 + \left(\frac{1+r}{\pi_j^e} - 1 \right) (1 - \tau_k), \tag{7}$$

where π_j^e is the conditional probability that an age $j-1$ individual with education e survives to age j. We assume that individuals can borrow up to an exogenous fixed limit ($a' \geq \underline{a}$) except for the last period of life.[11] Because the interest rate is adjusted by mortality rates, note that financial intermediaries break even when making loans (otherwise, intermediaries would go bankrupt). At age T we impose a non-negative constraint on savings ($a_{T+1} \geq 0$) capturing the fact that financial intermediaries are not willing to make loans to individuals that will be dead next period with probability 1. Note that there are no accidental bequests due to the presence of annuity markets. Nonetheless, from the perspective of households annuities are unfair because of the taxation of capital income.

3.4 The individual's problem

We use recursive language to describe the problem of an individual. To simplify the notation, we abstract from the fact that the education type of an individual determines his labor productivity process, mortality risk, and labor market shocks. Then, the state of an individual is

[10] This is done to avoid taking a stand on how accidental bequests are distributed across individuals. Note that the taxation of capital income makes the after tax return on annuities actuarially unfair.

[11] Keane and Wolpin (2001), in a structural model of educational attainment, find that borrowing constraints for youth are fairly tight. Keane and Wolpin (2001) also provide an overview of the literature on the importance of borrowing constraints on educational attainment and life-cycle consumption.

given by his age j, assets a, and earnings shock z, taste shock φ, and labor market status s. At the beginning of the period, prior to the labor supply decision, individuals can be in one of four labor market states s: with a job offer (e), without a job offer (n), unemployed (u), and retired (r). As discussed above, the only difference between being in the labor market states n and u is that in the latter individuals can collect unemployment insurance benefits.

Earnings can take one of several forms: labor earnings, unemployment insurance b_u, and pension benefits paid by the social security system b_s. Earnings are assumed to be taxed progressively according to a progressive tax schedule $T(y)$. The budget constraint is then given by

$$\begin{aligned} V_j(a,z,\varphi,s) &= \max_{\{c,h,a'\}} \{u(c,l,\varphi) + \beta\pi_{j+1}E[V_{j+1}(a',z',\varphi',s'(h,s))]\} \quad (8) \\ \text{s. t.} & \\ a' &= y(z,h,s) - T(y(z,h,s)) + R_j a - c\,(1+\tau^c), \\ a' &\geq \underline{a}, \end{aligned}$$

where the expectation operator is taken over productivity, taste, and labor market shocks. The labor market state next period, s', is a stochastic function of the current labor market state, s, and current working hours h, as described in Table 1. When individuals are retired, their labor market status is set permanently to r and earnings are given by the pension benefits. The only source of uncertainty faced by retired individuals is in their preference shock (φ) and mortality shocks.

3.5 Discussion on nonlinear earnings, fixed utility cost of participation, and labor supply decisions

The effects of nonlinear earnings on labor supply decisions can be illustrated with a simple static labor-supply problem. Consider an individual characterized by labor productivity z, non-labor income x, and a taste for leisure φ. The budget constraint is $c = zh^\theta + x$. The utility of working h hours is given by

$$W(h) \equiv (1-\varphi)\ln(z\,h^\theta + x) + \varphi\frac{(1-h)^{1-\sigma}}{1-\sigma} - F \quad (9)$$

Assuming that the individual chooses positive hours of work, the first-order condition (FOC) implies that the optimal amount of hours solves the following equation

$$\underbrace{\frac{(1-\varphi)z\theta h^{\theta-1}}{zh^\theta + x}}_{MB(h)} = \underbrace{\varphi(1-h)^{-\sigma}}_{MC(h)} \quad (10)$$

The MB and MC curves represent the marginal return and the marginal cost of working h hours measured in terms of utility. Figure 7 represents graphically the MB and MC curves for two different economies. Panel a) represents the curves for an economy with linear earnings ($\theta = 1$) and panel b) graphs the case for an economy with non-linear earnings ($\theta > 1$). In both economies, the MC curve is strictly increasing, convex, and satisfies $MC(0) = \varphi > 0$ and $\lim_{h\to 1} MC(h) = \infty$. The shape of the MB curve depends on whether the economy features linear earnings ($\theta = 1$) or not ($\theta > 1$).

Figure 7: The Effect of Non-linear Earnings on Labor Supply in a Static Model.

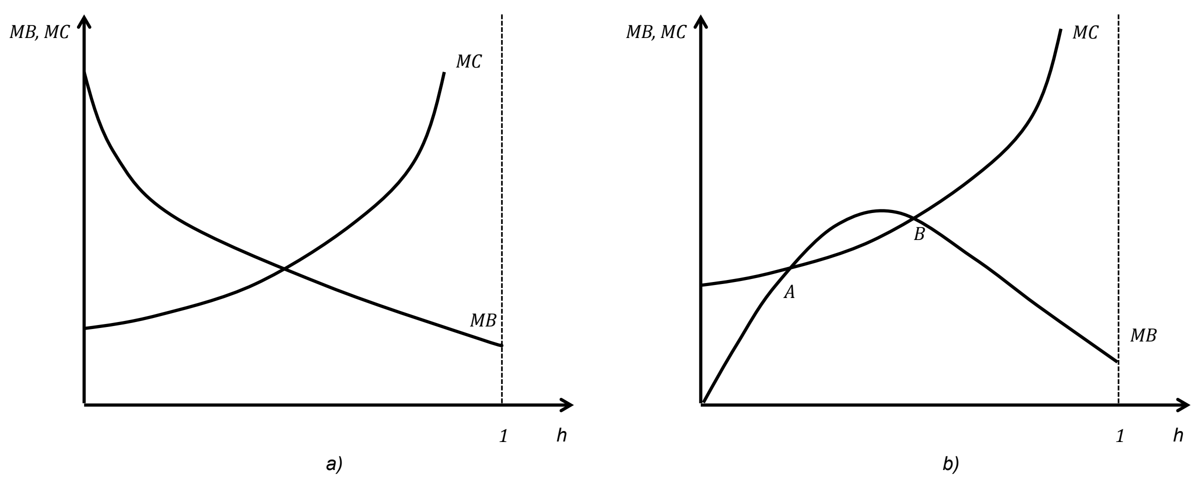

Case $\theta = 1$: Linear earnings. In this case, the wage rate is constant and the MB curve is strictly decreasing. As hours of work increase, consumption decreases and the marginal utility of consumption increases. As a result, the marginal benefit of working is strictly decreasing in hours of work and the MB and MC curves intersect at most once. When $\frac{1-\varphi}{x} < \varphi$, the MB and MC curves do not intersect and the optimal choice is not to work (the MC curve is above the MB curve for all $h \in [0, 1]$). This corresponds to a situation in which non-labor income (x) is high enough such that if the individual does not work ($h = 0$) his marginal utility of consumption ($\frac{1-\varphi}{x}$) is lower than the marginal utility of leisure (φ). When $\frac{1-\varphi}{x} > \varphi$, there is unique intersection and the optimal labor supply involves positive hours.

Case $\theta > 1$: Non-linear earnings. In this case, the MB curve is a strictly concave function of h, has an inverse-U shape with a maximum at $\hat{h} = \left(\frac{(\theta-1)x}{z}\right)^{\frac{1}{\theta}}$ and satisfies $MB(0) = 0$. The shape of the MB curve is due to the interaction of two opposing forces. The MB curve initially rises with hours (h) because the return to work measured in *consumption* units increases with hours of work when earnings are non-linear. On the other hand, the fact that the marginal utility of consumption decreases with the level of consumption is a force that tends to decrease the marginal return to work in terms of utility. When hours are low, the first effect dominates and the marginal benefit to work in terms of utility increases with hours. As hours of work increase, the second effect becomes relatively more important and the return to work becomes decreasing with hours. Note that the fact that the MB curve is initially increasing with hours implies that optimal hours of work (conditional on employment being optimal) will be bounded away from zero. To make further progress, it is convenient to consider two possibilities in terms of the overlap between the MB and MC curves. Depending on parameter values, the first possibility is that the two functions do not cross. In this case, $MC(h) > MB(h)$ for all h and it is optimal not to work. The second case is the one illustrated in Figure 7b in which the two functions cross for some $h \in (0, 1)$. In general, the two functions will cross twice (unless the MC and MB curves intersect at a tangent point, a situation which has measure zero in the parameter space). The first crossing point, labelled as A in Figure 7b, is a local minimum while the second one, labelled as B in Figure 7b, is a local maximum. Furthermore, since $MC(0) > MB(0)$, there is another local maximum at $h = 0$ (corner solution). Below we analyze labor supply decisions along the intensive and extensive margins.

Intensive margin. We now focus on how marginal changes in labor productivity (z), non-labor income (x), and the taste for leisure (φ) affect hours of work along the intensive margin (h_B). Assume that there is an individual, given by a triple of (φ, z, x) for which employment is optimal and denote by h_B the local maximum with positive hours in Panel b. Note that increases in z and/or decreases in x shift the MB curve up, but do not affect the MC curve (see equation 10). If follows that optimal hours along the intensive margin (h_B) increase with z and decrease with x. Moreover, since proportional changes in x and z do not affect the position of the MB curve, hours of work depend on the ratio of x/z, but not on the level of x and z. An increase in φ shifts the MC curve up and the MB curve down, thereby decreasing the optimal hours of work. In summary, the analysis implies that

$$\frac{\partial h_B}{\partial \varphi} \leq 0, \quad \frac{\partial h_B}{\partial x} \leq 0, \quad \frac{\partial h_B}{\partial z} \geq 0, \quad \frac{\partial h_B}{\partial (\frac{x}{z})} = 0.$$

Extensive margin. It is optimal to work when the difference between the utility of working h_B and the utility of non working is greater than (or equal to) the fixed cost of work. Using the expression for W in equation (9), it is optimal to work when $W(h_B) - W(0) \geq F$, which can be expressed as

$$W(h_B) - W(0) = \int_0^{h_B} W'(h)dh = \int_0^{h_B} [MB(h) - MC(h)]\, dh > F. \tag{11}$$

In terms of the graph in Panel b on Figure 7, it is optimal to work (i.e., h_B is a global maximum) when the difference in the areas under the MB and MC curves between 0 and h_b is greater than (or equal to) F. Note that on the graph shown in Figure 7b the difference between these two areas is negative, so that the optimal choice is to work zero hours even if fixed costs of work are zero ($F = 0$). Since for low hours of work (h close to zero) $MB(h) - MC(h) < 0$, the integral in expression (11) is negative for low values of h_B. It follows that individuals will only find optimal to work if h_B is bounded away from zero. This result is due to non-linear earnings and is reinforced by fixed costs of work.

Denote by $G(h_b) \equiv W(h_B) - W(0)$ the utility gain of working h_B hours, gross of the fixed cost of work F. Leibnitz's rule implies that

$$\frac{\partial G(h_B)}{\partial \varphi} = [MB(h_B(\varphi), \varphi) - MC(h_B(\varphi), \varphi)]\, \frac{\partial h_B}{\partial \varphi} + \int_0^{h_B} \left[\frac{\partial MB(h, \varphi)}{\partial \varphi} - \frac{\partial MC(h, \varphi)}{\partial \varphi}\right] dh < 0,$$

where we have used the fact that

$$MB(h_B(\varphi), \varphi) > MC(h_B(\varphi), \varphi), \quad \frac{\partial h_B}{\partial \varphi} < 0, \quad \frac{\partial MB(h, \varphi)}{\partial \varphi} < 0, \quad \frac{\partial MC(h, \varphi)}{\partial \varphi} > 0.$$

An increase in the disutility of work, reduces the gain from working h_B hours and makes it more likely that the individual will choose not to work.

Using a similar argument, it is easy to show that

$$\frac{\partial G(h_B)}{\partial x} < 0, \quad \frac{\partial G(h_B)}{\partial z} > 0.$$

Moreover, since (h_B, MC, MB) are independent of the ratio x/z, then

$$\frac{\partial G(h_B)}{\partial \left(\frac{x}{z}\right)} = 0.$$

Summing up, the gains from being employed are higher for individuals with high labor productivity (z) relative to their non-labor income (x) and for individuals with low taste for leisure (φ).

Frisch elasticity of labor supply: intensive margin. To obtain an expression for the Frisch elasticity of labor supply along the intensive margin, it is convenient to take logs on both sides of (10):

$$\ln(1 - h) = -\frac{1}{\sigma} \left\{ \ln(z\theta h^{\theta - 1}) - \ln(c) + \ln\left(\frac{1 - \varphi}{\varphi}\right) \right\}. \tag{12}$$

Differentiating with respect to the log of marginal earnings keeping constant consumption implies:

$$\frac{\partial \ln(1-h)}{\partial \ln(z\theta h^{\theta-1})} = -\frac{1}{\sigma}. \tag{13}$$

Using that the the hourly wage rate ($zh^{\theta}/h = zh^{\theta-1}$) is proportional to marginal earnings ($z\theta h^{\theta-1}$), it follows that

$$\frac{\partial \ln(1-h)}{\partial \ln(zh^{\theta-1})} = \frac{\partial \ln(1-h)}{\partial \ln(z\theta h^{\theta-1})} = -\frac{1}{\sigma}. \tag{14}$$

The Frisch elasticity of leisure (η_l) is equal to $-\frac{1}{\sigma}$. Since it does not depend on θ, it is equal to the expression obtained in a model with linear earnings ($\theta = 1$). Hence, nonlinear earnings do not affect the Frisch-elasticity of *leisure* along the intensive margin (see Aaronson and French (2009) for a related discussion.) The Frisch elasticity of *labor supply* along the intensive margin is then $\eta_h = \eta_l \frac{1-h}{h}$. The Frisch elasticity of labor varies across individuals. Individuals that work low hours relative to the mean in the population have a relatively high Frisch elasticity of labor supply along the intensive margin. The Frisch elasticity of aggregate labor supply along the intensive margin depends on the distribution of labor hours in the population.

Frisch elasticity of labor supply: extensive margin. We now analyze the factors affecting the Frisch elasticity of labor supply on the extensive margin and how these factors vary across the population. Assume that there is an initial situation in which some individuals find it optimal not to work. Consider the gains from work for individuals that choose not to work and have the same taste for leisure (φ). Revealed preferences imply that the gains from work in equation (11) are lower than F. Moreover, the gains from work are higher for individuals with low levels of consumption and high values of labor productivity, given that the MB curve shifts upwards with decreases in c and with increases in z. Now, assume that labor productivity increases by a factor of $1+\lambda$ for all individuals. Given a fixed consumption (c) level, the increase in labor productivity (z) shifts up proportionally the MB curve and does not affect the MC curve (see equation 10). Moreover, the shift in the MB curve is proportional to the change in z and inversely proportional to the initial (fixed) level of consumption (c), as shown in the expression below.

$$\Delta MB(h) = \frac{(1-\varphi)\theta h^{\theta-1}}{c}\Delta z = \frac{(1-\varphi)\theta h^{\theta-1}}{c}\lambda z. \tag{15}$$

Ceteris paribus, individuals with high labor productivity z and/or low consumption c are more likely to change their employment decision after labor productivity increases by a factor of λ for two reasons: First, these individuals faced higher initial gains from work and, second, their gains from work increase by a higher amount with the change in labor productivity. Similar arguments imply that, given c, individuals with low taste for leisure (φ) exhibit higher gains from work and these gains are more responsive to changes in labor productivity. It follows from this discussion that the compensated aggregate employment response to a change in productivity is determined by a host of factors: The preference parameters on leisure (σ, φ), the fixed cost of work (F), the extent of non-linear earnings (θ), and the joint distribution of individual characteristics over (φ, x, z). All of these objects determine the distribution of consumption and the distribution of gains from work in the population.

Summarizing, we have shown the following properties of labor supply decisions in our economy with non-linear earnings:

- Intensive margin: Hours of work along the intensive margin are bounded away from zero. They decrease when non-labor income (x) and taste for leisure (φ), and they are increasing in labor productivity (z). The Frisch-elasticity of leisure is determined by the curvature of the utility function on leisure (σ) and it is not affected by the assumption of non-linear earnings. The Frisch elasticity of labor supply along the intensive margin decreases with hours of work and is given by $\eta_h = \frac{1}{\sigma}(1-h)/h$.

- Extensive margin: Non-linear earnings penalize individuals working low hours because marginal earnings are close to zero when hours are low. Non-linear earnings encourages individuals to work either long hours or not at all. This tradeoff makes the extensive margin of labor supply decisions more prominent, even in the absence of fixed costs of work $F = 0$. Ceteris paribus, individuals that choose not to work are characterized by either a high value of non-labor income (x), taste for leisure (φ), or low productivity (z). The aggregate Frisch-elasticity of labor supply along the extensive margin is determined by preference parameters (φ, σ), technology parameters (θ) and the joint distribution of individuals characteristics (x, z, φ) in the population.

- Homotheticity: A doubling of labor earnings (z) and non-labor income (x), has no consequences for labor supply decisions along the extensive and intensive margins.

4 Calibration

We divide the parameters of the model economy in two groups. The first group includes parameters that are pinned down without simulating the model economy (such as, mortality rates and tax rates). The second group is composed of parameters that are calibrated by simulating the model economy.

4.1 Parameters calibrated without simulating the model economy

The model period is set to one quadrimester (4 months). The model economy is solved in partial equilibrium for a fixed interest rate. The quadrimesterly interest rate is chosen so that the implied annual rate of return on capital (net of depreciation) is 4%.

The intertemporal elasticity of substitution of leisure. In the calibrated baseline economy we choose $\sigma = 2.0$ which implies an intertemporal elasticity of leisure of 0.5. We also report the results for an alternatively calibrated economy with $\sigma = 3$ (an intertemporal elasticity of leisure of 1/3).

Non-linear earnings. The hourly wage in our theory satisfies

$$w_h(z,h) = \frac{zh^{\theta}}{h} = zh^{\theta-1}. \tag{16}$$

Note that the elasticity of the wage rate to a change in hours of work is given by $(\theta - 1)$. As we discussed in Section 2.3, in an empirical study Aaronson and French (2004) estimate this elasticity to be around 0.40. This estimate implies that a full time (40 hours a week) worker earns an hourly wage 20% higher than a part time (20 hours) worker. We thus set $\theta = 1.4$.

Tax rates, social security, and unemployment insurance benefits. The tax rate on consumption τ_c is set at 0.075 as in McDaniel (2007). Following Domeij and Heathcote (2004), taxes on capital income are set to $\tau_k = 0.40$. The social security tax rate is set to $\tau_{ss} = 0.153$, and the cap $\overline{y}$ on social security taxation is fixed at 2.47 of average earnings in the economy ($\overline{y}$). Taxes on earnings (non-capital income) are set according to a progressive tax schedule with 5 tax brackets. Following French (2005), we parameterize the tax schedule using data on US federal and state income taxes for household heads (with the standard deduction) in 1990 in the state of Rhode Island, which is a fairly representative state in terms of its income tax system. The tax brackets, are defined by the following thresholds expressed as multiples of the average earnings ($\overline{y}$) in the economy: $0.10, 0.16, 0.63, 1.93$. The corresponding tax rates for the five tax brackets are $0\%, 13.2\%, 17.9\%, 32.9\%, 36.9\%$.

As discussed before, for computational simplicity we follow a vast literature in modeling a stylized representation of the US social security system. For each education level, we assume two different values for average lifetime earnings (one value for each of the two possible fixed effects on labor productivity). Then the average earnings for each of the four productivity types are computed according to the benefit formula of the US social security system (the US social security benefit formula is a function of the Primary Insurance Amount and has two bend points at 0.2 and 1.24 of the average earnings in the economy).

The unemployment insurance benefit is set to 40% of the potential earnings of the unemployed worker, where potential earnings are computed assuming a work week of 40% of available time.

Mortality rates. The mortality risk for college and non-college individuals is taken from Bhattacharya and Lakdawalla (2006).

Exogenous job separation rates into unemployment. Based on the empirical analysis in Section 2.1 we set the exogenous separation rates into unemployment (δ) in a quadrimester to 2.4% for non-college and and 1.16% for college.

4.2 Parameters calibrated by simulating the model economy

Rather than simulating the model to calibrate all the parameters at once, we find it convenient to partition the parameters in two subgroups and to follow an iterative procedure with two nested loops. The parameters in each of the two subgroups are calibrated in two separate loops in order to diminish the dimensionality of the calibration. In the inner loop, for fixed values of the parameters calibrated in the outer loop, we calibrate the parameters determining the stochastic process on taste for leisure, the job finding rate, and the fixed utility cost of work. Given the parameters obtained in the inner loop, the outer loop calibrates the parameters determining the labor productivity process, discount rate, average earnings in the economy (used to define tax schedules), average lifetime earnings for the four fixed productivity types (used to compute pension benefits for the two fixed productivity types in each of the two education groups) and the exogenous borrowing limit.

4.2.1 Inner loop

The inner loop pins down the parameters determining the heterogeneity in the taste for leisure (φ), the fixed utility cost of work (F), and the job finding rate (p). Recall that taste for leisure (φ) depends on an individual fixed effect (φ_i) and a stochastic shock (φ_t). The fixed effect is drawn from a normal distribution, $\varphi_i \sim N(\mu_\varphi, \sigma_\varphi^2)$. The stochastic shock ($\varphi_t$) represents a deviation from the mean value and follows a first order auto-regressive process with persistent value ρ_φ and innovation η_φ (see equation (3)). The innovation is drawn from a normal distribution with variance $\sigma_{\eta_\varphi}^2$. The initial value of φ_t (at age 25) is drawn from the invariant distribution so that the cross-sectional distribution of taste shocks does not vary over the life cycle.[12] Hence, we need to pin down, for each education group, six parameters: four

[12] The continuous stochastic process in the model for the taste for leisure is approximated with a Tauchen procedure, with nine possible values (three for the fixed effects and three for the persistent taste shock).

values determining the heterogeneity in taste for leisure ($\mu_\varphi, \sigma^2_\varphi, \rho_\varphi, \sigma^2_{\eta_\varphi}$) plus the fixed utility cost of leisure F and the job finding rate p. The values of these parameters are obtained by minimizing the distance (square of the sum of deviations) between the model and the SIPP data on the following statistics:

1. Average employment rates in a quadrimester for four age groups (25-34, 35-44,45-54, and 55-61), as computed in Section 2.1 from the SIPP.

2. The probability of entering a non-employment spell in a quadrimester for four age groups (25-34, 35-44,45-54, and 55-61), as computed in Section 2.1 from the SIPP.

3. The fraction of all non-employment spells lasting one, two, three, or more than three quadrimesters, as computed in Section 2.1 from the SIPP.

4. The coefficient of variation of lifetime hours for workers between the ages of 35 and 44, as computed in Section 2.2 from the PSID: 0.35 for non-college and 0.23 for college individuals, respectively.

5. Average mean hours of work in a quadrimester among prime-age males. Following Osuna and Ríos-Rull (2003) and Prescott (2004), the time endowment is set at 5200 hours a year (100 hours per week). Using data from the 1990 SIPP on prime-age males with positive hours of work, 35 to 50 years old, we obtain that non-college and college individuals work, on average, about 41.2% and 43.5% of their time endowments in a quadrimester.

4.2.2 Outer loop

The outer loop pins down the parameters governing the labor productivity process, discount rate, average earnings in the economy (used to define tax schedules), average lifetime earnings for four productivity types (in order to compute social security benefits), and the exogenous borrowing limit. The labor productivity process is calibrated, through an iterative procedure that ensures that the model is consistent with data moments on hour hourly wages over the life cycle. Since the SIPP is longitudinally fairly short to allow us to estimate the stochastic process for wages, we use the PSID for this purpose. In calibrating a *quadrimesterly* stochastic process on labor productivity, one difficulty arises from the fact that the PSID only reports earnings and hours of work at an annual frequency. Moreover, we need to consider that the data only report wages for individuals that work. To deal with these problems, we follow an iterative procedure:

1. Estimate a wage profile and wage process for college and non-college workers on annual data from the PSID data. In particular, we specify the following annual process for log hourly wages:

$$\ln(\widetilde{w}_h)_{ij} = \widetilde{\mathbf{x}}_{\mathbf{j}}\widetilde{\kappa} + \widetilde{\alpha}_i + \widetilde{u}_j + \widetilde{\lambda}_j, \tag{17}$$

where $\ln(\widetilde{w}_h)_{ij}$ represents the observed annual log hourly wage of individual i at age j in the PSID data, $\widetilde{\mathbf{x}}_{\mathbf{j}}$ is a quartic polynomial in age, $\widetilde{\kappa}$ is a vector of coefficients, $\widetilde{\alpha}_i \sim N\left(-\frac{\sigma^2_{\widetilde{\alpha}}}{2}, \sigma^2_{\widetilde{\alpha}}\right)$ is a fixed effect determined at birth, $\widetilde{\lambda}_j \sim N\left(-\frac{\sigma^2_{\widetilde{\lambda}}}{2}, \sigma^2_{\widetilde{\lambda}_j}\right)$ is an idiosyncratic transitory shock which is interpreted as measurement error, and $\widetilde{u}_j$ follows a first-order autoregressive process:[13]

$$\widetilde{u}_j = \rho_{\widetilde{u}}\widetilde{u}_{j-1} + \widetilde{\eta}_{u_j}, \quad \widetilde{\eta}_{u_j} \sim N\left(-\frac{\sigma^2_{\widetilde{\eta}_u}}{2}, \sigma^2_{\widetilde{\eta}_u}\right), \quad \widetilde{u}_0 = 0. \tag{18}$$

2. Feed a *quadrimesterly* labor productivity process $\Omega = (\kappa, \rho_u, \sigma^2_\alpha, \sigma^2_{\eta_u})$ into the model economy as specified in Section 3.1.[14]

3. Simulate the model economy to obtain *quadrimesterly* data on employment, hours of work, and earnings.

4. Aggregate the *quadrimesterly* data to an annual period.

5. Estimate an annual hourly wage profile and hourly wage stochastic process for college and non-college workers in the model generated data.

6. Feed a new *quadrimesterly* labor productivity process (go back to step 2), until the "same" annual wage profile and stochastic wage process is obtained in the model and in the data.

Following Kaplan and Violante (2010), the discount factor β is chosen to match an asset to income ratio of 2.5. This is the wealth to income ratio when the top 5% of households in the wealth distribution are excluded from the Survey of Consumer Finances (SCF). The reason for excluding the richest households in computing an aggregate wealth to income ratio is that the PSID under-samples the top of the wealth distribution. The borrowing limit ($\underline{a}$) is pinned down so that the model is consistent with the consumption growth rate between ages

[13] We follow Meghir and Pistaferri (2004) and Low et al. (2010) in interpreting the purely iid component in residual log wages (or annual earnings) as measurement error. Meghir and Pistaferri (2011) provide an overview of the literature and the available specifications of the stochastic processes of wages (and earnings).

[14] A Tauchen procedure is used to approximate the stochastic process with a finite number of realizations of the shock. We use two values for the fixed effect and 20 values for the persistent shock.

25 to 55 of about 25% (see Fernández-Villaverde and Krueger (2007)). We assume a common value of β and $\underline{a}$ for the two education groups since we target aggregate data, across the two education groups, on the asset to income ratio and consumption growth over the life cycle.

4.3 Calibration results

4.3.1 Outer loop

Our calibration procedure implies a value for the discount factor of $\beta = 0.983$ and a borrowing constraint equivalent to 40% of average annual earnings in the economy. The latter estimate is similar to the one obtained by Kaplan (2012). As argued in Kaplan (2012), if borrowing were not allowed consumption growth predicted by the theory would be counterfactually high. The calibration implies a capital to output ratio of 2.6 and a mean consumption growth over the life cycle of 24%, which are close to the targets of 2.5 and 25% respectively. Table 2 reports the values of the parameters driving the stochastic process on labor productivity in the baseline economy. These parameters were determined using the iterative procedure described above (see section 4.2.2). The stochastic process on wages estimated on the model data and on the PSID data deliver similar estimates. Moreover, Figure 8 shows that the model mimics well the life cycle profile of wages.

Figure 8: Deterministic life-cycle wage profile, PSID and model.

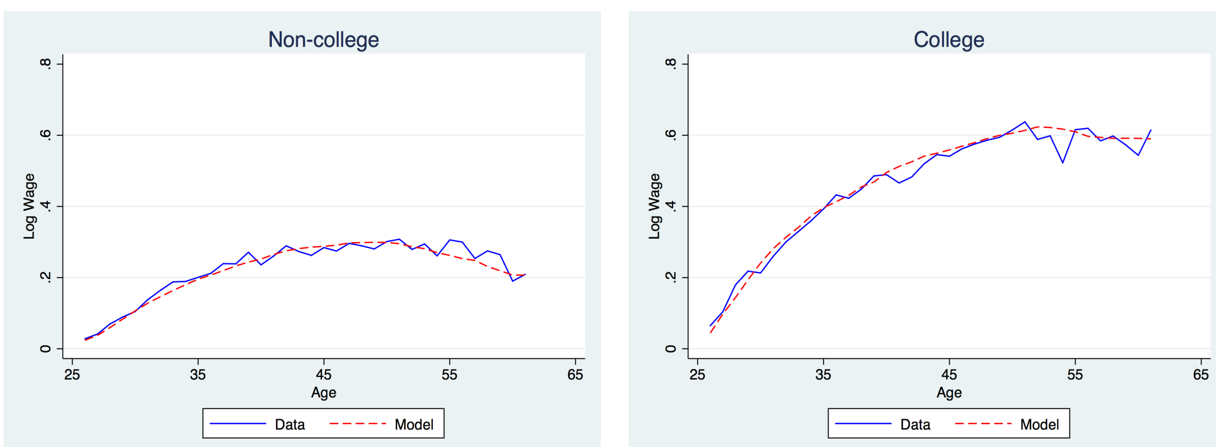

4.3.2 Inner loop

Table 3 reports the parameters obtained in the inner loop. The calibration implies that individuals from both education groups face a job finding rate below one (0.69 and 0.59 for

Table 2: Calibration of the stochastic process for wages.

Non-College				
Parameter	Value	Moment	Data	Model
σ^2_{α}	0.094	$\sigma^2_{\widetilde{\alpha}}$	0.094	0.097
ρ_u	0.949	$\rho_{\widetilde{u}}$	0.941	0.940
$\sigma^2_{\eta_u}$	0.014	$\sigma^2_{\widetilde{\eta}_u}$	0.019	0.019
College				
Parameter	Value	Moment	Data	Model
σ^2_{α}	0.065	$\sigma^2_{\widetilde{\alpha}}$	0.075	0.074
ρ_u	0.982	$\rho_{\widetilde{u}}$	0.968	0.977
$\sigma^2_{\eta_u}$	0.016	$\sigma^2_{\widetilde{\eta}_u}$	0.019	0.020

non-college and college) and substantial fixed (utility) costs of work. To make sense of the importance of the fixed costs of work, it is convenient to rewrite the utility function as follows:

$$\begin{aligned} u(c_t, l_t) &= (1-\varphi)\ln c_t + \varphi \frac{l_t^{1-\sigma}}{1-\sigma} - I_{(l_t<1)}F, \\ &= (1-\varphi)\ln\left(\frac{c_t}{1+\hat{F}}\right) + \varphi\frac{l_t^{1-\sigma}}{1-\sigma} \\ \text{where } \hat{F} &\equiv \exp\left(\frac{F}{1-\varphi}\right) - 1 \text{ if } l_t < 1, \text{ and equal to 0 otherwise.} \end{aligned}$$

The calibration implies the utility cost of work is equivalent to losses in consumption of 8.5% and 10.5% for non-college and college individuals, respectively. In Section 6, we argue that these numbers are consistent with the evidence from Aguiar and Hurst (2013) on work-related expenditures. The calibration implies substantial heterogeneity in the preference shocks for the non-college and college individuals, with a coefficient of variation of around 0.20 for both education groups.[15]

4.3.3 Other moments targeted in the calibration

We now show how the baseline model economy matches the moments targeted in the calibration. We emphasize that, except for the age-profile of labor productivity, all the parameters

[15] The patterns in the preference heterogeneity are consistent with the analysis in Heathcote et al. (2014) who, using different identifying restrictions, reach similar conclusions.

Table 3: Calibration of parameters in the inner loop.

Variable	Parameter	Non-college	College
Prob. of finding a job	p	0.69	0.59
Mean of fixed effect of taste for leisure	μ_{φ_i}	0.60	0.55
Std. dev. of fixed effect of taste for leisure	σ_{φ_i}	0.12	0.11
Pers. of transitory shock of taste for leisure	ρ_{φ_t}	0.73	0.21
Std. dev. of transitory shock of taste for leisure	σ_{η_φ}	0.006	0.008
Fixed cost of work (% consumption loss)	$\hat{F}$	0.085	0.105

in the model do not vary with age. In particular, our calibration assumes that the job separation and job finding rates, the fixed cost of work, and the distribution of preference shocks are constant over the life cycle. Hence, the life-cycle patterns shown below are the result of the model features driving life-cycle behavior: a finite time horizon, asset accumulation for retirement, precautionary savings, and the age-profile of wages.

Employment rates. The baseline economy is consistent with the fact that the age-profile of employment is roughly flat early in the life cycle for both education groups (see Figure 9). The model tends to over-predict employment rates, though not by much. The baseline economy captures the fact that for non-college individuals employment starts declining at age 45 and that this decline accelerates substantially after age 55. Moreover, consistent with the data, the decline in the employment rate for college individuals starts at an older age than for non-college. Overall, the model captures quite well, qualitatively and quantitatively, the decline in employment rates late in the life cycle.

Hours worked along the intensive margin and lifetime labor supply. The model economy matches the targets of 0.412 and 0.435 for the average hours of work for prime-aged college and non-college individuals (aged 35-50), respectively. Matching these targets requires that the weight on leisure in the utility function be higher for non-college than for college individuals (0.60 versus 0.55).

The model is also consistent with the heterogeneity in lifetime labor supply in the PSID data. The calibration targeted the coefficient of variation of lifetime hours of work for the 35-44 age group, which takes the value of 0.35 for non-college and 0.23 for college individuals, respectively (see Section 2.2). To match this fact, the calibration requires substantial heterogeneity in the taste for leisure for both education groups, with a bigger variance for

Figure 9: Employment rate, 1990 SIPP and model.

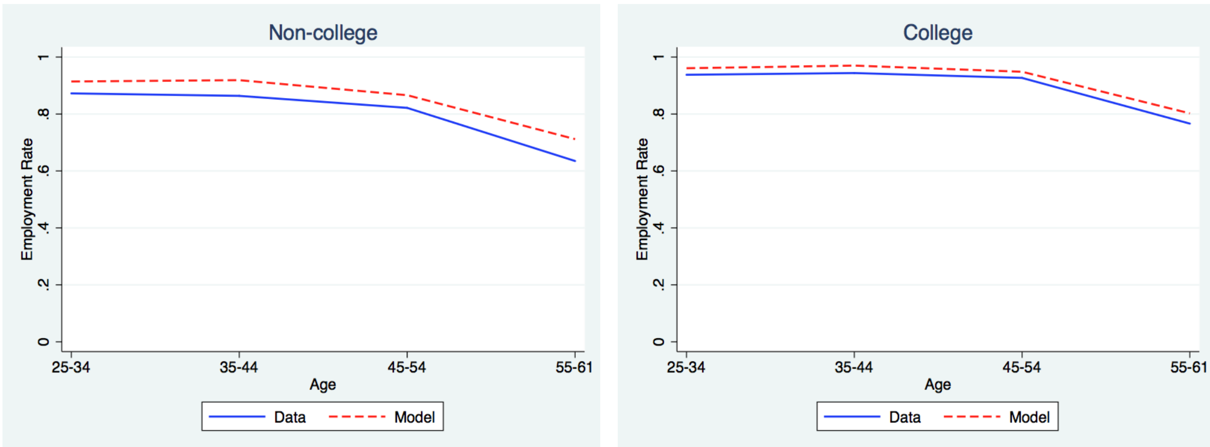

non-college than college individuals.

Entry into non-employment spells. The calibration targets the quadrimesterly hazard rates into non-employment over the life cycle. Note that the transition rates from employment into non-employment provide information on the variation of labor supply decisions along the extensive margin. Hence, this target is important for disciplining the predictions of the theory for how the aggregate employment rate responds to temporary wage shocks. Figure 10 shows that the model matches quite well the hazard rates into non-employment. The baseline economy matches well the fact that the hazard rates have a U-shape over the life cycle and are higher for non-college than college individuals. We emphasize that the life cycle profile for hazard rates are not due to variation in preference shocks over the life cycle. As we shall discuss later, allowing for (positive) borrowing and modeling preference heterogeneity is important for the decline in hazard rates early in the life cycle. The accumulation of assets over the life cycle and the increase in the variance of labor productivity with age is important for the increase in the transitions from employment into non-employment late in the life cycle.

Non-employment spells. The length of non-employment spells provides a measure of the "volatility" of labor supply decisions along the extensive margin. If people that quit jobs go back to work the next period, then labor supply decisions along the extensive margin are quite "volatile." We thus think it is important that the baseline economy is consistent with evidence on entry into non-employment as well as with the duration of non-employment spells. An important achievement of our theory is that it is consistent with evidence on both the incidence and the volatility of non-employment. Figure 11 shows that the model matches fairly well, qualitatively and quantitatively, the statistics on the length of non-employment

Figure 10: Entry rate into non-employment, 1990 SIPP and model.

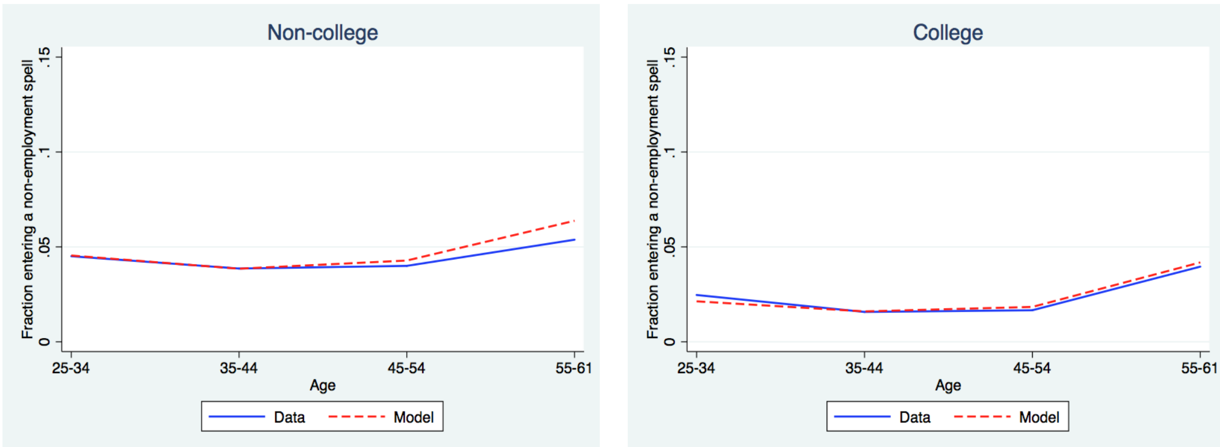

Notes: The entry rate into non-employment is measured as the fraction of those employed in quadrimester t that are non-employed in quadrimester $t+1$.

spells documented on SIPP data. Non-employment spells are divided in four groups depending on their length: 1, 2, 3, or 4+ quadrimesters. For both education groups, in the SIPP and model data, more than 50% of all non-employment spells last one period and the fraction of non-employment spells lasting 4 or more periods is about 20%.

Figure 11: Non-employment spells, 1990 SIPP and model.

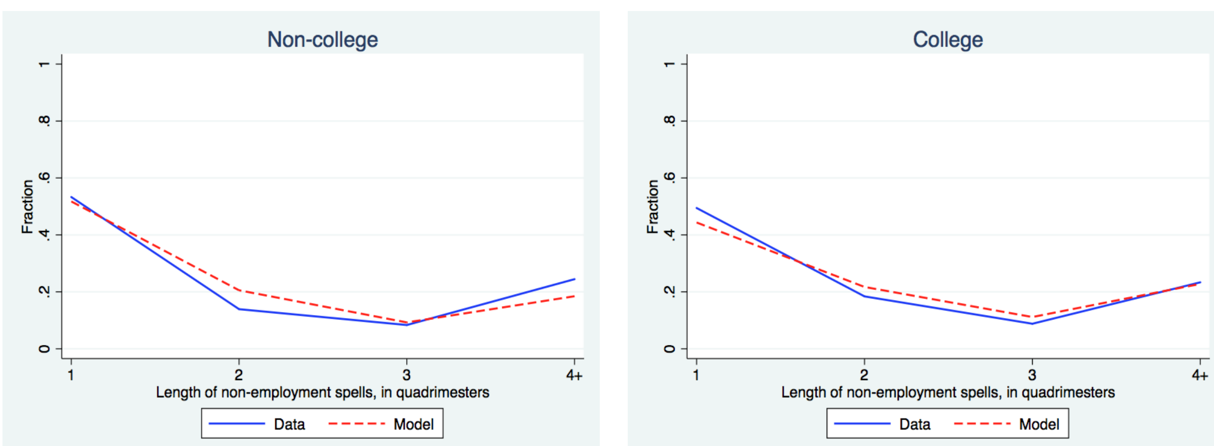

4.3.4 Moments not targeted in the calibration

The model also performs very well in dimensions not directly targeted in the calibration procedure.

Hours worked. The calibration targeted the employment rate for four age groups and the mean hours worked for employed prime-age males (35 to 50 years old). Figure 12 shows that the model matches the age-profile of average hours worked in the population (including individuals with zero hours worked) quite well. Figure 13 shows that the overall distribution of hours worked in a quadrimester in the baseline model economy is similar to the distribution observed in the 1990 SIPP.

Figure 12: Mean hours, 1990 SIPP and model.

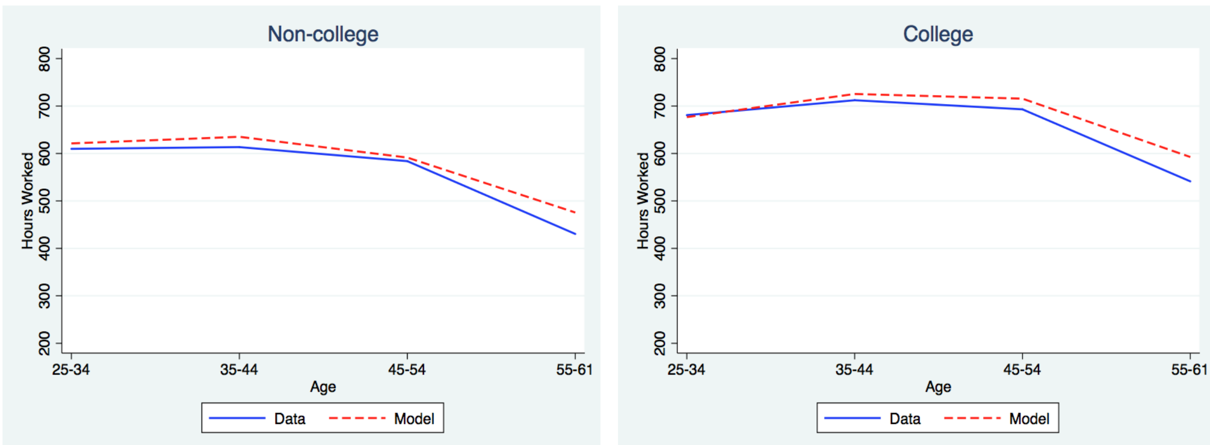

Figure 13: Histograms of hours worked in a quadrimester, 1990 SIPP and benchmark model.

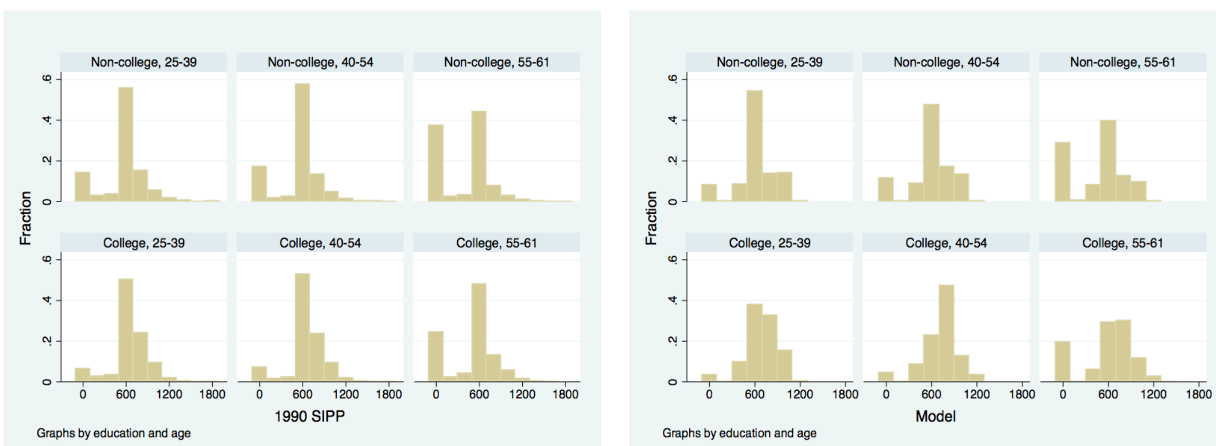

Non-employment spells by age. The calibration targeted statistics on the duration distribution of *all* non-employment spells. Figure 14 shows that the model captures quite well how the duration distribution of non-employment spells varies over the life cycle in the SIPP data. When individuals are young, both for non-college and college individuals, about 60% of the non-employment spells last for only one period. As individuals age, this fraction decreases and is about 40% for the age group 55-61. On the other hand, for both educational categories, the fraction of non-employment spells that are of length equal to four or more periods grows over the life cycle reaching a value close to 40% for the age group 55-61. The model matches these facts closely.

Figure 14: Non-employment spells, 1990 SIPP and model, by age.

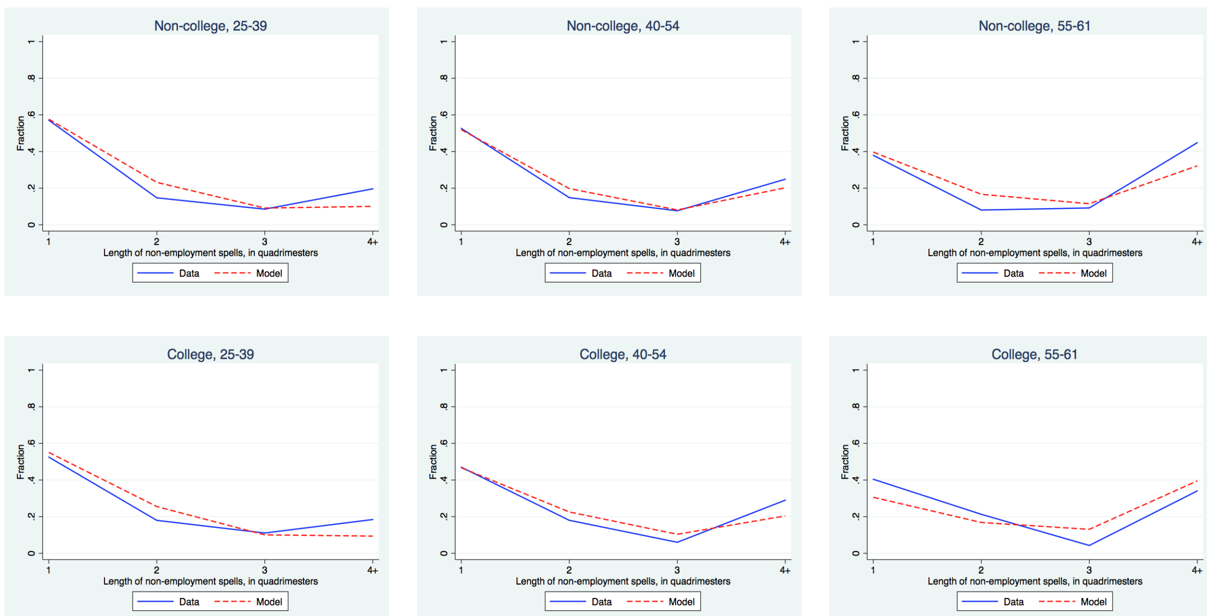

Cross-sectional heterogeneity in hours and lifetime labor supply by age. Consistent with the SIPP data, the model predicts that the coefficient of variation in hours of work is large, grows with age, and at the end of the life cycle reaches a value above 0.6 for non-college and of about 0.4 for college individuals (see Figure 15). Recall that our calibration procedure implies that the variance of taste for leisure is constant over the life cycle. Hence, the increase in the variance of working hours with age is due to an increase in the variances of labor productivity and wealth over the life cycle. For all age groups, both in the model and in the data, heterogeneity in labor supply is larger for non-college than college individuals.

The calibration targeted the coefficient of variation of lifetime hours for the age group 35-44. The baseline economy matches remarkably closely how inequality in lifetime labor supply (the coefficient of variation over the sum of hours worked during a 10 year period) increases over the life cycle and that the increase in the inequality of lifetime hours over the life cycle is larger for non-college than college individuals. Indeed, as seen on Figure 16, quantitatively the

Figure 15: Coefficient of variation in hours, PSID and model.

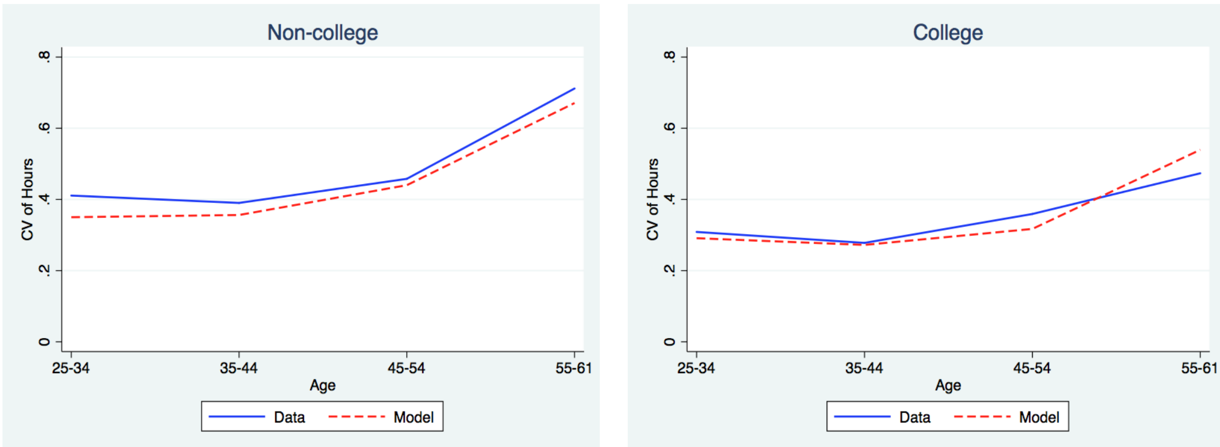

Notes: Let h_{ij} denote the annual hours worked for individual i who is j years old. The coefficient of variation in hours across all individuals of age j is $CV^j(h_{ij})$. The figure reports the average $CV^j(h_{ij})$ over a ten-year period (e.g., ages 35-44), $E_j[CV^j(h_{ij})]$.

model matches the age-profile of lifetime hours inequality very well for both education groups. Since the variance of preference shocks is constant over the life cycle, the increase in lifetime-hours inequality is explained by two mechanisms embedded in the model: life cycle behavior and incomplete markets. Individuals build precautionary savings early in the life cycle to ensure against idiosyncratic risk. As they age and accumulate assets, the heterogeneity in asset holdings translate into more persistent differences in labor supply across individuals.

Consumption inequality over the life cycle. We have shown that the baseline economy is consistent with life cycle data data on US inequality in wages, labor supply, and lifetime labor supply. Huggett et al. (2011) and Kaplan (2012) advocate the view that a life-cycle theory of inequality and incomplete markets should be broadly consistent with evidence on consumption inequality over the life cycle. It is therefore interesting to compare the baseline economy's implications for the rise in consumption dispersion over the life cycle with the patterns found in U.S. data. The model economy implies a variance of log consumption of 0.17 at age 25 and of 0.30 at age 60. Hence, consumption inequality rises by 13 log points between the ages of 25 and 60. These implications are consistent with the evidence on consumption inequality reviewed in Huggett et al. (2011). In particular, Aguiar and Hurst (2013) document that the increase in the variance of consumption is about 12 log points when consumption is measured as total nondurable expenditures, with an initial value at age 25 of 0.15. The model's implications are remarkably close to these estimates. The Gini index of consumption in the baseline economy

Figure 16: Coefficient of variation in lifetime hours, PSID and model.

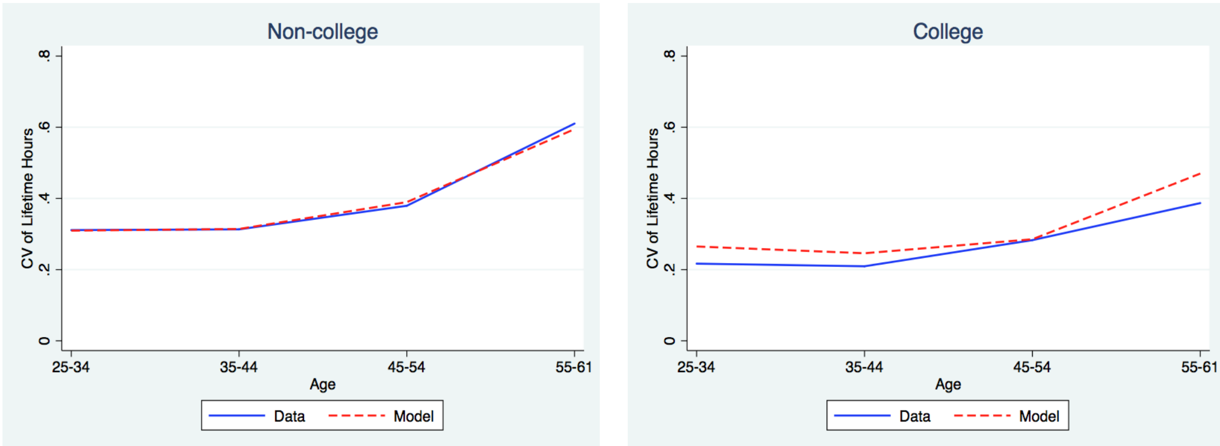

Notes: For each ten year period (e.g., between the ages of 35 and 44), let h_{ij} denote the annual hours worked for individual i who is j years old and let $\tilde{h}_i = (\sum_{j=35}^{44} h_{ij})$ denote the individual i's lifetime hours during this period. The coefficient of variation in lifetime hours across all i individuals is $CV(\tilde{h}_i)$.

is 0.28, which is again close to the value of 0.26 reported by Krueger and Perri (2006) for the US economy during the 1990-2000 period.

5 Aggregate labor supply responses

The analysis so far has shown that the baseline economy is able to match well facts on labor supply along the extensive and intensive margins at the individual level. Therefore, the model economy is an appropriate tool for studying aggregate labor supply responses to changes in the economic environment. The model is used to explicitly aggregate up each individual's response, as well as to analyze how responses vary across subgroups in the population. Moreover, in a series of experiments we shut down one by one various modeling assumptions in order to evaluate their importance for understanding aggregate labor supply responses.

5.1 Response to a temporary wage change

We start by simulating in the baseline economy the aggregate labor supply response to a one-period (quadrimester) unanticipated wage change of 2%. In principle, one may think that the wealth effect of such a change is negligible so that the change in the aggregate labor supply provides an estimate of the Frisch elasticity of aggregate labor supply. However, this reasoning is not firmly grounded in economic theory as with incomplete markets and borrowing constraints the Frisch elasticity of labor supply is an ill-defined concept.[16] First, with borrowing constraints the Frisch elasticity does not exist because agents cannot hold the marginal utility of consumption constant as wages fluctuate. Second, with heterogeneity and incomplete markets a change in the wage rate will affect the marginal utility of wealth differently for different types of people. Hence, our results in this section should be viewed as describing aggregate labor supply responses to a temporary wage change rather than the Frisch elasticity of labor supply.[17,18]

The aggregate elasticity of labor supply to a temporary wage change can be decomposed in terms of the elasticities along the intensive and extensive margins. The change in the average hours of work among employed individuals is used to compute the elasticity along the intensive margin. The elasticity along the extensive margin is given by the elasticity of the employment rate.[19] The results from these computations are reported in Table 4. We find that the elasticity of aggregate labor supply with respect to a temporary wage change in the baseline economy is 1.75. Restricting attention to labor supply changes along the intensive margin decreases the response from 1.75 to 0.67. Hence, the extensive margin accounts for about 62% of the aggregate labor supply response to a temporary wage change.

The elasticity of labor supply to a temporary wage change varies importantly across education and age groups. It is higher for high school than for college educated individuals (1.96 versus 1.35, respectively). Labor supply responses have a strong life cycle pattern: The age-profile of the elasticity is U-shaped for both education categories. For high-school individuals, the elasticity of labor supply takes a value of 2.01 for the 25-34 age group and decreases to 1.62 for the 35-44 age group. Afterwards, it increases to 1.90 for the 45-54 age group and to 2.74 for the 55-61 age group. The variation of the elasticity of labor supply for college individuals over the life cycle exhibits a similar U-shaped pattern. Such a pattern is reminiscent of the fact that over the business cycle hours of work and employment fluctuate much more for young and old individuals than for the middle-aged, as discussed in Gomme et al. (2004) and Jaimovich and Siu (2009).

The variation of labor supply responses across education and age groups in the baseline economy is almost all due to changes along the extensive margin. To fix ideas, let us focus

[16] See Keane (2011) for a discussion.

[17] We thank an anonymous referee for pointing this out.

[18] We note that in our baseline economy about 93.5% of all individuals experience a less than 0.2% change in their consumption in response to a 2% wage change implying that the variation caused by a wage change in the marginal utility of wealth across agents in the baseline economy is small.

[19] Aggregate hours of work H can be expressed as $H = e \times h$, where e denotes the employment rate and h mean working hours. Taking logs and differentiating with respect to the log-wage gives an expression for the aggregate elasticity in terms of the extensive an intensive margin elasticities: $\eta_H = \eta_e + \eta_h$.

Table 4: Aggregate labor supply elasticity to a temporary wage change.

	ALL			HS	COL
Elasticities:	Total	Intensive	Extensive	Total	Total
1. Baseline Economy	1.75	0.67	1.08	1.96	1.35
2. Low Productivity Risk	1.42	0.69	0.73	1.60	1.09
3. No Labor Market Frictions	1.85	0.68	1.17	2.04	1.49
4. No Credit	1.51	0.67	0.84	1.64	1.27
5. Linear Earnings + No Fixed Cost	0.85	0.85	0.0	0.86	0.81
6. No Fixed Cost	1.27	0.70	0.57	1.39	1.04
7. Linear Earnings	0.85	0.85	0.0	0.86	0.82
8. No Preference Heterogeneity	1.29	0.71	0.58	1.37	1.13

on high school individuals. The intensive margin elasticity is almost flat over the life cycle: It decreases monotonically from 0.69 at age 25-34 to 0.66 at age 55-61. On the contrary, the extensive margin elasticity has a U-shape over the life cycle: It starts at 1.33 for the age group 25-34, decreases to 0.95 at age 35-44 and then rises to 1.23 and to 2.08 at ages 45-54 and 55-61, respectively (see Table 5). Hence, the U-shaped age profile of the labor supply response to temporary wage change is thus driven by the variation of labor supply responses along the extensive margin. A similar conclusion applies to college educated individuals. Furthermore, the extensive margin is crucial for understanding differences in the elasticity of labor supply across education groups. In short, the intensive margin is quantitatively important, but is not a factor in understanding the variation in the elasticity of aggregate labor supply to a temporary wage change across age and education groups. All the variation in labor supply responses across these groups is due to the extensive margin.

5.1.1 Evaluating the importance of modeling assumptions for the response to a temporary wage change

The baseline economy models (i) productivity risk, (ii) labor market risk, (iii) credit markets, (iv) fixed cost of work, (v) non-linear earnings, and (vi) preference heterogeneity. We now conduct various experiments to isolate the importance of each of these modeling assumptions for the aggregate labor supply response to a temporary wage change. In each of these experiments we shut down one model assumption and keep the rest of the parameters of the baseline economy fixed. We then evaluate the labor supply change to a one period unanticipated wage change of 2% in each of the modified model economies. The results from these experiments are summarized below.

1. **Productivity risk**. To isolate the role of labor productivity risk, we assume that the standard deviation of the innovation to the autoregressive process on labor productivity is 100 times smaller. As a result, we obtain an economy in which individuals essentially do not face labor productivity risk (apart from the initial fixed effect). The rest of the parameters are kept equal to the ones in the baseline economy. Table 4, row 2, reports the elasticity of labor supply to a temporary wage change for the modified-baseline economy with no labor productivity risk. The elimination of productivity risk decreases the aggregate elasticity of labor supply from a value of 1.75 in the baseline economy to 1.42. This decrease is entirely due to a lower labor supply response at the extensive margin (from 1.08 to 0.73) and is mostly due to old individuals. When productivity risk is shut down, for both education groups, the elasticity along the extensive margin for individuals aged 55-61 decreases sharply: It decreases from 2.08 to 0.46 for non-college and from 1.42 to 0.39 for college individuals (see Table 5). This result is explained through the effect of productivity risk on life-cycle savings in the baseline economy. Since individuals have stronger incentives to save in order to self-insure against productivity risk, old individuals in the baseline economy are richer than in the economy with no risk and their employment decisions become more responsive to temporary wage changes.

2. **Labor market risk**. To assess the importance of labor market risk, we consider an economy in which the probability of finding a job is set to one ($p = 1$), keeping all other parameters from the baseline economy constant. The results are reported in row 3 of Table 4. When labor market risk is shut down the elasticity of aggregate labor supply increases from 1.75 in the baseline economy to 1.85. Interestingly, this result differs from the previous experiment, where we found that shutting down labor productivity risk led to a *decrease* in the aggregate labor supply response. In understanding why labor market risk and productivity risk have opposite effects on labor supply decisions, it is important to keep in mind that labor market risk affects labor supply decisions through two channels. First, it encourages precautionary savings which, as discussed

Table 5: Extensive margin elasticity to a temporary wage change, by age and education.

Panel A: High School	$25-34$	$35-44$	$45-54$	$55-61$
1. Baseline Economy	1.33	0.95	1.23	2.08
2. Low Productivity Risk	1.27	0.89	0.82	0.46
3. No Labor Market Frictions	1.09	0.94	1.45	2.68
4. No Credit	0.52	0.72	1.21	2.14
5. Linear Earnings + No Fixed Cost	0.00	0.00	0.00	0.00
6. No Fixed Cost	0.67	0.50	0.75	0.97
7. Linear Earnings	0.00	0.00	0.00	0.05
8. No Preference Heterogeneity	0.30	0.34	0.62	2.28

Panel B: College	$25-34$	$35-44$	$45-54$	$55-61$
1. Baseline Economy	0.64	0.45	0.58	1.42
2. Low Productivity Risk	0.60	0.40	0.27	0.39
3. No Labor Market Frictions	0.47	0.56	0.72	2.09
4. No Credit	0.31	0.42	0.59	1.46
5. Linear Earnings + No Fixed Cost	0.00	0.00	0.00	0.00
6. No Fixed Cost	0.28	0.29	0.34	0.60
7. Linear Earnings	0.00	0.00	0.00	0.04
8. No Preference Heterogeneity	0.27	0.26	0.31	1.18

before, makes employment decisions of old individuals much more responsive to temporary wage changes. Now, the strength of this effect is much weaker in the case of labor market risk than productivity risk because in our baseline model economy labor market shocks (job finding rates) are not persistent and thereby have a weaker effect on precautionary savings. Second, by making it difficult to find a job, labor market risk makes it more costly for individuals to take periods off work and the extensive margin becomes less responsive to temporary wage changes. We find that the second effect quantitatively prevails, especially late in the life cycle. As a result, as indicated in Table 5, the labor-supply elasticity of high school individuals aged 55-61 along the extensive margin *increases* from 2.08 to 2.68 when labor market risk is shut down (recall that this statistic *decreases* to 0.46 in the experiment eliminating labor productivity risk). For college individuals, this statistic *increases* from 1.42 in the baseline economy to 2.09 when labor market risk is eliminated (while it *decreases* to 0.39 in the absence of productivity risk).

3. **Credit markets**. In the baseline economy individuals can borrow up to about 40% of average annual earnings in the economy. Shutting down credit markets, decreases the elasticity of aggregate labor supply from 1.75 in the baseline economy to 1.51 (see row 4 in Table 4). The impact of credit markets on labor supply responses is almost all due to changes along the extensive margin. These effects vary over the life cycle in important ways. Focusing on high school individuals, the elimination of credit markets *decreases* the labor supply response along the extensive margin of individuals aged 25-34 from 1.33 to 0.52, while it *increases* the extensive margin elasticity of individuals aged 55-61 from 2.74 to 2.81 (see Table 5). Credit markets also have opposite effects on the labor supply responses of young and old individuals with college education. Since young individuals hold few assets, the absence of credit makes it more costly for young individuals to take periods off work, diminishing their labor supply responses along the extensive margin. On the other hand, the absence of credit encourages individuals to build precautionary savings making their employment decisions more responsive to temporary wage shocks when old.

4. **Non-linear earnings, fixed costs of work, and extensive margin responses**. We now consider labor supply responses in three different economies. First, we consider an economy with no fixed costs of work ($\hat{F} = 0$) and linear earnings ($\theta = 1$). Since in this economy the extensive margin is essentially inoperative, the elasticity of aggregate labor supply decreases from 1.75 to 0.84 relative to the baseline economy (see row 5 in Table 4).

The decrease in the aggregate labor supply response is all explained by a lower response along the extensive margin (from 1.08 to 0).[20] Hence, in an economy with no fixed costs of work and linear earnings the elasticity of labor supply along the extensive margin is zero. To evaluate separately the role of the fixed cost of work and non-linear earnings for the observed extensive margin responses, we evaluate employment responses when only one of these two model features is active. If we add the calibrated fixed cost of work to the economy with linear earnings, the elasticity along the extensive margin is essentially zero (see row 7 on Table 4), suggesting that the calibrated fixed cost of work is too small for the extensive margin to respond to a small temporary wage change.[21] If, instead, we add non-linear earnings to an economy with no fixed costs of work the extensive margin elasticity is 0.57, which is a substantial response but still half of the overall response of 1.08 (compare rows 1 and 6 in Table 4). The last observations underscores that the calibrated fixed costs of work are important for employment responses in the presence of non-linear earnings.

5. **Preference heterogeneity**. To evaluate how heterogeneity in preferences affects labor supply responses, we consider an economy in which the variance of the preference shocks is set to zero. When preference heterogeneity is shut down in the baseline economy, the elasticity of aggregate labor supply to a temporary wage change decreases from 1.75 to 1.29 (compare rows 1 and 8 in Table 4)). We find that eliminating preference heterogeneity reduces the aggregate employment elasticity by about a half (from 1.08 to 0.58), with this reduction affecting all education and age groups (see Table 5).[22] These effects are particularly large for individuals aged 25-34: The extensive margin elasticity falls by more than a fourth for high school individuals (from 1.33 to 0.30) and by more than a half for college individuals (from 0.64 to 0.27). Preference heterogeneity leads to a higher employment elasticity because it allows for "lazy" types who value leisure strongly. These individuals are much less likely to work than "non-lazy" individuals and, conditional on working, they are much more likely to enter a non-employment spell. As a result, "lazy" types exhibit a high employment elasticity to temporary wage changes.

[20] The intensive margin elasticity is higher in the economy with an inactive extensive margin due to the fact that working hours are lower in this economy. This can be understood by recalling that $\eta^h = \frac{1-h}{h}\eta^l$, where η^h and η^l denote the elasticity of hours and leisure, and h denotes hours of work. Since the elasticity of leisure is fixed at $1/\sigma$, a decrease in h mechanically increases the elasticity of hours.

[21] We will further discuss this issue in section 5.4.2.

[22] The only exception is for high school individuals aged 55-61, whose employment elasticity increases from 2.08 to 2.28 without preference heterogeneity.

Summary of findings. We find that the elasticity of aggregate labor supply to a temporary wage change in the baseline economy is 1.75, with the extensive margin accounting for about 62% of the aggregate labor supply response. Moreover the labor supply elasticity is much higher for high school individuals than college educated individuals (1.96 versus 1.35). Through a series of experiments aimed at assessing the importance of various modeling assumptions, we find that the intensive margin, although quantitatively important, matters little for understanding the variation in the labor supply elasticities across all the model economies. The impact of modeling assumptions on labor supply responses is almost entirely driven by their effects on decisions along the extensive margin. When we simulate employment responses to a small wage change in an economy with linear earnings, we find that the calibrated fixed costs of work are sufficiently small that, on their own, they do not affect employment decisions along the extensive margin. However, fixed costs of work matter importantly in the presence of non-linear earnings, accounting for half of the employment response in our baseline economy. This result underscores that there is a complementarity between fixed costs of work and non-linear earnings that enhances the aggregate labor supply response to temporary wages.

The extensive margin is also crucial for understanding how the labor supply responses to temporary wage change varies across age and education groups. The age profile of the labor supply elasticity in the baseline economy is U-shaped. Modeling credit (or initial assets), preference heterogeneity, and the extensive margin are important for generating the decline in the elasticity early in the life cycle. If any of these features are shut down in the baseline economy, the elasticity of labor supply does not decrease with age when individuals are young. Modeling productivity risk and the extensive margin (non-linear earnings) are important for accounting for the increase in the labor supply response to a temporary wage change late in the life cycle. Preference heterogeneity increases aggregate labor supply responses, and this effect is strongest for young individuals with high school education. Finally, labor market frictions have a large negative effect on the labor supply response of old individuals.

Our model abstracts from human capital accumulation. It is interesting to point out that in a model with human capital accumulation, as in Imai and Keane (2004) and Keane (2011), the labor supply response to a temporary wage change increases over the life cycle. Young individuals do not respond much to temporary wage shocks because the opportunity cost of work is not only the current wage but also the return to human capital accumulation. As individuals age, the returns to human capital accumulation decline and individuals respond more to temporary wage changes. Our model focuses on different, but complementary, mechanisms which have similar implications as the ones in Imai and Keane (2004). For instance,

Table 6: Aggregate labor supply elasticity to a permanent-compensated wage change.

	Total	Extensive	Intensive
Baseline Economy	0.44	-0.11	0.55
Linear Earnings	0.56	0.00	0.56
No Preference Heterogeneity	0.49	-0.07	0.56

the need to build precautionary savings makes young individuals less responsive to temporary wage changes while as they age and accumulate assets they become more responsive.

5.2 Response to a permanent compensated wage change

We now evaluate the aggregate labor supply response to a permanent compensated wage change, which provides an estimate of the Hicks elasticity of labor supply. This is done by simulating a permanent wage decrease of 10% together with a lump sum transfer that effectively compensates the decrease in earnings due to the lower wage rate. Hence, the wage change affects the first order condition for working hours and earnings are only affected to the extent that individuals change hours of work. To put it differently, the compensated wage change does not directly affects earnings but only indirectly through its impact on hours of work.

We find that the Hicks elasticity of aggregate labor supply in the baseline economy is 0.44.[23] Interestingly, labor supply responses along the intensive and extensive margins have different signs. While the intensive margin elasticity is 0.55, the extensive margin elasticity is -0.11. Hours along the intensive margin decrease because of the substitution effect associated with the lower wage rate. On the other hand, the decrease in the (compensated) wage discourages savings over the life cycle which, in turn, leads to an important increase of the employment rate at old age. This effect accounts for the aggregate negative elasticity of labor supply along the extensive margin. In particular, the Hicks elasticity along the extensive margin

[23] This number is consistent with the evidence summarized in Chetty et al. (2011). Keane (2011) finds similar to our paper effects to a compensated permanent tax increase in a model with human capital accumulation. Interestingly, however, Keane (2011) shows that in a model with human capital accumulation it is important to distinguish between the short-run and long-run effects of a compensated permanent tax change. Since a permanent increase in taxes decreases the incentives to accumulate human capital, the effects of such a tax reform on labor supply grow over time and are thus higher in the long-run than in the short-run.

for individuals aged 55-61 is -0.42 and -0.35 for non-college and college individuals. The employment elasticity for younger age groups is much smaller (less than one fifth the value for the oldest age group).

We also evaluated the Hicks elasticity in two separate experiments: In the first experiment we shut down non-linear earnings (relative to the baseline economy), whereas in the second experiment we shut down preference heterogeneity. We find that in the case with linear earnings the aggregate Hicks elasticity is 0.56, which is higher than the 0.44 estimated for the baseline economy. The Hicks elasticity is now higher than in the baseline economy because the extensive margin elasticity is zero instead of negative. Shutting down preference heterogeneity also increases the aggregate Hicks elasticity from 0.44 in the baseline economy to 0.49. The aggregate Hicks elasticity is now slightly higher because the extensive margin elasticity decreases from -0.11 in the baseline economy to -0.07 in the absence of preference heterogeneity. Hence, unlike the findings for the response to a temporary wage change, preference heterogeneity does not matter much for the Hicks elasticity of labor supply.

5.3 Intertemporal substitution and labor supply late in the life cycle

There are substantial cross-country differences in labor supply late in the life cycle (age 50+). In many countries the social security provisions impose explicit and implicit taxes on the labor earnings of individuals that are close to the normal retirement age, encouraging these individuals to reduce their labor supply and withdraw from the labor market before the normal retirement age (Gruber and Wise (1999)). The cross-country heterogeneity in labor supply and social security provisions present us with the opportunity to test the predictions of our theory. In Erosa et al. (2012) we use a related framework with non-linear earnings to model in detail the variation in the social security, disability insurance, and taxation institutions across European countries and the United States. We find that the model economy accounts well for the observed cross-country differences in labor supply late in the life cycle, indicating that the intertemporal labor supply responses in our framework plausible.

5.4 Iceland tax experiment

In a recent paper Chetty et al. (2011) evaluate whether state-of-the-art macro models featuring indivisible labor are consistent with modern quasi-experimental micro evidence. They simulate the Rogerson and Wallenius (2009) model in order to assess the implications of the theory for labor supply responses to temporary and permanent tax (wage) changes. They find that the model is consistent with the micro estimates of the Hicks elasticity (steady-state compensated tax changes) of labor supply but that it is inconsistent with the micro evidence on the Frisch elasticity of labor supply. Since our model economy is essentially an extension of the Rogerson and Wallenius (2009) model to a setting with incomplete markets and heterogeneous agents, it is interesting to evaluate the predictions of our theory relative to the evidence discussed in Chetty et al. (2011).

Chetty et al. (2011) test the Rogerson and Wallenius (2009) model with evidence from Bianchi et al. (2001) on labor supply responses to a tax policy change in Iceland during the 1987. This policy change implied that during the year 1987 income was untaxed in Iceland. Chetty et al. (2011) argued that this tax holiday induced an unanticipated temporary wage variation that provided a natural experiment for identifying Frisch elaticities. They simulated a one year unanticipated tax relief in the Rogerson and Wallenius (2009) (henceforth RW) framework and concluded that the RW implied a labor supply response much larger (about a factor of five) than the estimated response in the data.

Our model economy has not been calibrated to Iceland (neither the RW model). Nonetheless, the findings of Chetty et al. (2011) raise the concern that our theory may be grossly at odds with evidence from the Icelandic Tax holiday. To evaluate this possibility, we simulate the Icelandic tax holiday in our baseline economy. In 1987, Iceland moved from a system under which taxes were paid on the previous year's income to a pay-as-you-earn system. The transition to the new tax system implied that income during 1987 was never taxed since the tax base in 1987 was income earned in 1986 and the tax base in 1988 was income earned in 1988.[24] The average tax rate was 14.5% in 1986, 0% in 1987, and 8.0% in 1988. In order to mimic the tax reform in Iceland, we simulate in our baseline economy a one year (three model periods) reduction in the tax rate of 14.5 percentage points, followed then by a permanent decrease of 6.5 percentage points in the average tax rate.

Before proceeding to the results, we want to clarify the objective of this experiment. First, we emphasize that the Icelandic tax holiday does not provide evidence on Frisch elasticities of labor supply or the aggregate labor supply response to a transitory wage change. One reason is due to the fact that Iceland was undergoing substantial policy reforms in 1987 which were likely to affect future growth prospects and, in particular, affect the marginal utility of wealth of the agents in the economy. Moreover, as we have stated before, the 1987 tax holiday was accompanied by some permanent tax changes after 1987. The Icelandic tax holiday combines a temporary tax relief with a permanent decrease in taxes. Hence, the results should be

[24]The tax change was unanticipated by households since the announcement of the policy change was made in late 1986, see Bianchi et al. (2001).

interpreted with care. Our goal is to test if our theory grossly overestimates the labor supply responses estimated in the Icelandic data, as argued by Chetty et al. (2011) for the RW model.

We find that the aggregate elasticity of labor supply implied by the Icelandic tax holiday experiment is 0.64. Remarkably, this elasticity result is not much different from the ones estimated by Bianchi et al. (2001) in the Icelandic micro data who found an aggregate labor supply elasticity of 0.84 for male workers.[25]

6 Sensitivity analysis and discussion

In this section we study labor supply responses in two additional model economies. The first economy is identical to the baseline economy but re-calibrated under the assumption of linear earnings. The second economy is also identical to the baseline economy but it is re-calibrated under the assumption that preferences on leisure feature $\sigma = 3$. The calibration strategy and the targets are the same as those used in the parameterization of the baseline economy.

6.1 Comparing model economies

While we find that all economies do a decent job in matching the calibration targets, the economy with linear earnings is the one that has the worst fit of the calibration targets.[26] Instead of comparing a long list of statistics across model economies, we now highlight the most noticeable differences between the calibrated model economies.

The most striking difference across model economies is that the baseline economy implies much lower fixed costs of work than the other two economies (see Table 7). In particular, for non-college, the economy with linear earnings and the economy with $\sigma = 3$ require a fixed cost of work equivalent to a consumption loss of 36% and 24%, respectively, which are much higher than the 8.5% estimated in the baseline economy. These findings should not be surprising: Individuals are less willing to take periods off work when, ceteris paribus, earnings

[25] This value comes from adding the extensive and intensive margin elasticities. They report an extensive margin elasticity of 0.58 in Table 4 in Bianchi et al. (2001). The intensive margin elasticity of 0.26 is obtained using data from Table 6. In our simulations the response along the extensive margin is 0.23, which is substantially smaller than the 0.58 value reported by Bianchi et al. (2001). However, we should note that these authors used data on number of weeks worked to measure the extensive margin response. Since our model period is a quadrimester, our higher response along the intensive margin is probably capturing an extensive margin response in the data studied by Bianchi et al. (2001). We thus think that it is more appropriate to focus our analysis on aggregate labor supply responses.

[26] For high school individuals, the average deviation of the model statistics relative to the calibration targets is 0.116 in the baseline economy, 0.196 in the economy with linear earnings, and 0.137 in the economy with $\sigma = 3$. For college individuals, these statistics are 0.116 in the baseline economy, 0.161 in the economy with linear earnings, and 0.109 in the economy with $\sigma = 3$.

Table 7: Calibration of parameters in the inner loop, three economies.

NON-COLLEGE	Parameter	Baseline $\sigma = 2$	Linear earnings $\sigma = 2$	$\sigma = 3$
Prob. of finding a job	p	0.69	0.68	0.70
Mean of fixed effect of taste for leisure	μ_{φ_i}	0.60	0.53	0.48
Std. dev. of fixed effect of taste for leisure	σ_{φ_i}	0.12	0.18	0.18
Pers. of trans. shock of taste for leisure	ρ_{φ_t}	0.73	0.85	0.75
Std. dev. of trans. shock of taste for leisure	σ_{η_φ}	0.006	0.007	0.009
Fixed cost of work (% consumption loss)	$\hat{F}$	0.085	0.36	0.24
COLLEGE	Parameter	Baseline $\sigma = 2$	Linear earnings $\sigma = 2$	$\sigma = 3$
Prob. of finding a job	p	0.59	0.67	0.62
Mean of fixed effect of taste for leisure	μ_{φ_i}	0.55	0.48	0.44
Std. dev. of fixed effect of taste for leisure	σ_{φ_i}	0.11	0.12	0.15
Pers. of trans. shock of taste for leisure	ρ_{φ_t}	0.21	0.98	0.67
Std. dev. of trans. shock of taste for leisure	σ_{η_φ}	0.008	0.01	0.007
Fixed cost of work (% consumption loss)	$\hat{F}$	0.105	0.37	0.21

are a linear function of hours of work or when preferences exhibit higher curvature on leisure than in the baseline economy. As a result, the calibrations of these two economies require much higher fixed costs of work than in the baseline economy in order to match the age-profile of employment. We now discuss some evidence from Aguiar and Hurst (2013) that suggests that the estimated fixed costs in the two newly calibrated economies are implausibly large.

In a recent study, Aguiar and Hurst (2013) argue that expenditures on (i) non-durable transportation, (ii) food away from home, and (iii) clothing/personal care should be viewed as inputs into market labor supply.[27] They labelled the expenditures on items in these three categories as "work-related" expenditures and show that they are crucial for understanding the hump-shape of consumption and the increase in the cross-sectional dispersion of consumption over the life cycle. Using consumption data from the Consumption Expenditure Survey (CEX), they estimate that expenditures in these three categories amount to about 28% of household consumption expenditures at age 43-45. However, as discussed by Aguiar and Hurst (2013), the actual amount of work-related expenditures incurred by households is only a fraction of this number. This is because the data on spending on transportation, food away from home,

[27]This view follows a long tradition in labor economies. See for instance Cogan (2001) and the references in Aguiar and Hurst (2013).

and clothing/personal care likely bundles spending on goods that are associated with work as well as goods that represent non-work spending. For instance, part of the expenses in transportation expenditures reflect the need to commute to work as well as travel for other (leisure) purposes. Because the expenditure data set does not distinguish costs due to work travel from costs due to non-work travel, they use time diaries from the pooled 2003-2005 American Time Use Survey (ATUS) to gauge the relative importance of each. The detailed categories of the ATUS allows them to identify traveling to and from work separately from time spent traveling for other reasons. They find that for those who work, work-related travel represents roughly 30% of all travel expenses. Moreover, Aguiar and Hurst (2005) find evidence that expenditures on food away from home decline by 30% after retirement. Hence, assuming that 30% of clothing/personal care represent costs of work, we obtain an estimate of fixed costs of work of 0.30×0.28, which roughly amounts to about 9% of consumption expenditures. This is similar to the fixed costs of work estimated in the baseline economy: 8.5% for non-college and 10.5% for college. An alternative back of the envelope calculation on the importance of work-related expenditures can be made as follows: Aguiar and Hurst (2005) find that work-related expenses drop by 40% after retirement. Assuming that this drop is entirely explained by the decrease in labor supply we obtain an estimate of fixed costs of work of about 12% (0.28×0.40). We conclude that the fixed cost of work of 36% (24%) estimated in the calibration of the economy with linear earnings (the economy with $\sigma = 3$) appears to be too high relative to the evidence on work-related expenditures studied by Aguiar and Hurst (2005, 2013).

Another important difference across the calibrated model economies is that the baseline economy features much lower preference heterogeneity among non-college individuals than the other two economies. The standard deviation of the fixed effect on taste for leisure in the baseline economy is 0.12 whereas it is 0.18 in the other two economies. Preference heterogeneity has consequences for the distribution of hours of work, inequality in lifetime hours, and the variance of consumption. We find that the economy with linear earnings performs worse than the other economies in several dimensions. Figure 17 shows the histograms of hours worked for non-college individuals (i) in the baseline model economy, (ii) in the linear re-calibrated model economy, and (iii) in the re-calibrated model economy with $\sigma = 3$. As we pointed out earlier, the distributions of hours worked in the data and in the baseline model are roughly similar. However, the distribution of hours worked in the linear re-calibrated model exhibits some counterfactual patterns. The variance of hours is large relative to the data: There is a large fraction of people in the 200-400 hours interval and there is a very large fraction

Figure 17: Histograms of hours worked, non-college, three model economies.

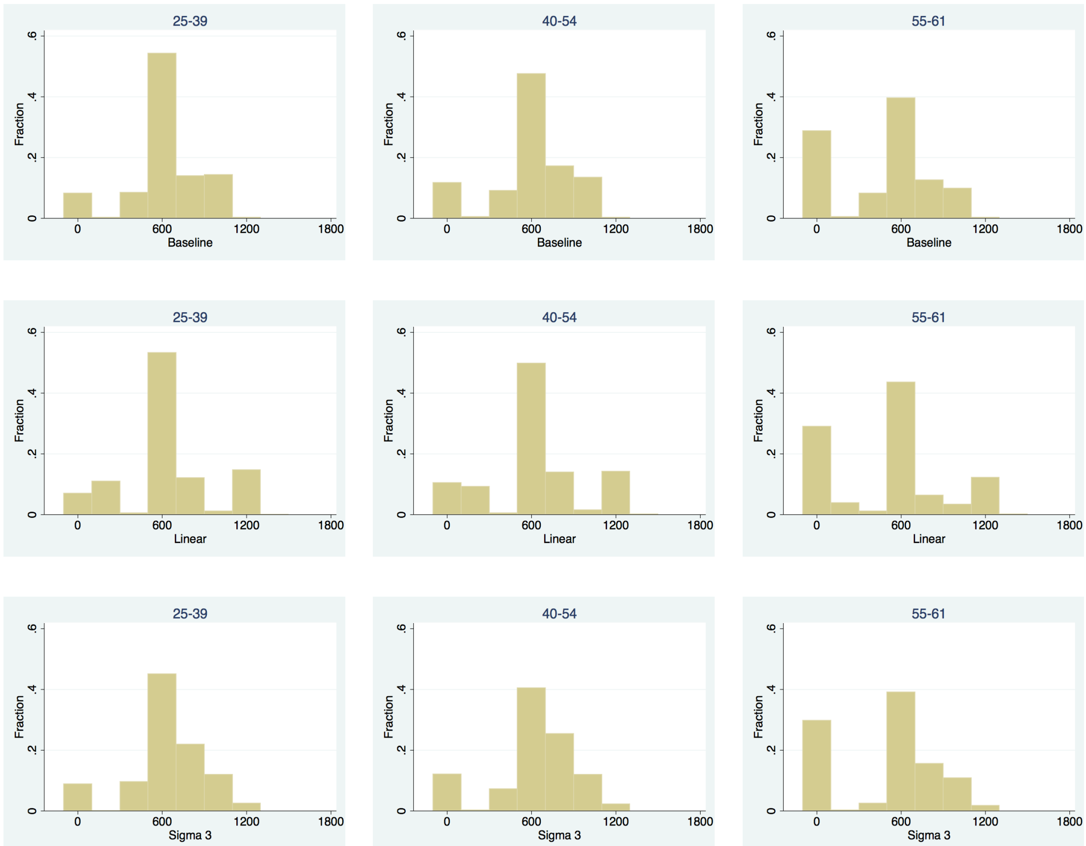

Notes: The figure shows the histograms of quadrimesterly hours for the non-college, age 25-61, in three model economies: the baseline economy, the recalibrated linear earnings economy, and the recalibrated economy with $\sigma = 3$.

of workers working long hours (in the 1200-1400 hours interval). Obviously, this is a direct consequence of the high preference heterogeneity featured by the linear re-calibrated model economy. Surprisingly, even though the economy with $\sigma = 3$ has about the same amount of preference heterogeneity as the linear re-calibrated economy, the dispersion in hours in the former economy is not as high as in the latter. The reason is that the economy with $\sigma = 3$ features non-linear earnings and less curvature on leisure, which reduces the amount of workers working either short hours or long hours.

An additional implication of the high preference heterogeneity estimated in the economy with linear earnings is that this economy has a high variation in lifetime labor supply relatively to the data. In effect, this economy has a coefficient of variation in lifetime labor supply of 0.47, which is much higher than the 0.35 value in the data. This statistic is equal to 0.36 and 0.37 in the baseline economy and in the economy with $\sigma = 3$, respectively. Hence, the results suggest that non-linear earnings are important for matching the evidence on inequality in lifetime labor supply. Furthermore, we examine how the variation in preference heterogeneity impacts on consumption inequality across economies. In the baseline economy, consistent with the empirical estimates reviewed by Huggett et al. (2011), the variance of log consumption among non-college individuals is 0.14 at age 25 and rises to 0.27 at age 60. In the economy with linear earnings, these statistics − with values of 0.20 and 0.37 − appear to be higher relative to the empirical evidence.

6.2 Labor supply responses across model economies

We now compare labor supply responses across the three calibrated model economies. We also consider an additional economy that is identical to the baseline economy but that only differs in that the curvature parameter on leisure is fixed at $\sigma = 3$. This economy is simulated to isolate only the impact of a change in σ, keeping all other parameters in the economy fixed at their values in the baseline economy.

Temporary wage change. We start by evaluating the response to a one period (quadrimester) wage change of 2%. We find that the economy with linear earnings has a much lower aggregate response than the baseline economy, with an elasticity of aggregate labor supply of 1.05 versus 1.75 in the baseline economy (see Table 8). The lower response in the re-calibrated economy with linear earnings is explained by the extensive margin. In this economy the elasticity of aggregate labor supply along the extensive margin is less than a half the one in the baseline economy (0.44 versus 1.08). Recall the in section 5.1 we have shown that if we introduce linear earnings in the baseline economy, without recalibrating the model, the extensive margin response was essentially zero. We have then calibrated an economy with linear earnings and found that this economy features fixed costs of work that are about four times higher than in the baseline economy. We draw two key lessons from simulating a temporary wage shock in this re-calibrated linear earnings economy: First, fixed cost of work may allow for significant (bounded away from zero) extensive margin responses if they are large enough. Second, non-linear earnings *amplify* the effect of fixed costs of work on extensive margin responses to a temporary wage change. The baseline economy (with non-linear earnings) has an extensive margin response to a 2% wage increase that is 2.5 times higher (e.g. the ratio between 1.08 and 0.44 is about 2.5) than the re-calibrated economy with linear earnings even though fixed cost of work are about a factor of four smaller in the former economy.

The re-calibrated economy with $\sigma = 3$ implies a response to a temporary wage change of 1.42, which is smaller but not much smaller than the 1.75 value in the baseline economy. The

Table 8: Aggregate labor supply elasticity to a temporary wage change: sensitivity.

	Baseline $\sigma = 2$	Linear earnings $\sigma = 2$	Baseline $\sigma = 3$	Baseline $\sigma = 2$ with $\sigma = 3$
Elasticities:				
Aggregate	1.75	1.05	1.42	1.22
Non-college	1.96	1.15	1.60	1.31
College	1.35	0.86	1.06	1.05
Extensive margin	1.08	0.44	0.94	0.59
Non-college	1.29	0.54	1.12	0.67
College	0.69	0.24	0.56	0.44
Intensive margin	0.67	0.61	0.48	0.63
Non-college	0.67	0.61	0.48	0.64
College	0.66	0.62	0.50	0.61

response along the extensive margin is 0.94, which is not far from the 1.08 obtained in the baseline economy. Moreover, the variation of the extensive margin elasticity across age and education groups follow similar patters across the two model economies: (i) the age profile of the extensive margin elasticity is U-shaped, with the oldest age group (55 to 61 years of age) featuring the largest responses to a temporary wage change (see Table 9); (ii) for all ages, the employment decisions of non-college individuals are more responsive to a small temporary wage change than the ones of college individuals.

To isolate the role of the preference parameter σ we simulated an economy identical to the baseline economy but with $\sigma = 3$ (instead of $\sigma = 2$). We find that the extensive margin elasticity of aggregate labor supply decreases by almost a half (from 1.08 to 0.59). Intuitively, individuals care more about smoothing leisure over time when $\sigma = 3$ than when $\sigma = 2$, which accounts for the large reduction on the employment response. However, the re-calibrated economy with $\sigma = 3$ exhibits an employment elasticity of 0.94, much larger than the 0.59 value for the modified baseline economy with $\sigma = 3$. This is because the former economy has higher fixed costs of work and more preference heterogeneity than the baseline economy, which undo most of the reduction of the extensive margin elasticity associated to an increase in σ. These results underscore that the extensive margin elasticity is not determined by σ but jointly by all parameters in the model economy (such as wage process, fixed costs of work, preference heterogeneity).

Table 9: Extensive margin elasticity to a temporary wage change, by age and education.

Panel A: High School	$25-34$	$35-44$	$45-54$	$55-61$
Baseline $\sigma = 2$	1.33	0.95	1.23	2.08
Linear Earnings $\sigma = 2$	0.36	0.23	0.51	1.68
Baseline $\sigma = 3$	1.11	0.80	1.04	2.00
Baseline $\sigma = 2$ parameters but $\sigma = 3$	0.46	0.43	0.77	1.42
Panel B: College	$25-34$	$35-44$	$45-54$	$55-61$
Baseline $\sigma = 2$	0.64	0.45	0.58	1.42
Linear Earnings $\sigma = 2$	0.08	0.03	0.14	1.08
Baseline $\sigma = 3$	0.52	0.32	0.45	1.26
Baseline $\sigma = 2$ parameters but $\sigma = 3$	0.32	0.28	0.39	0.97

Table 10: Aggregate labor supply elasticity to a permanent-compensated wage change.

	Baseline $\sigma = 2$	Linear earnings $\sigma = 2$	Baseline $\sigma = 3$
Elasticities:			
Aggregate	0.44	0.41	0.33
Non-college	0.43	0.40	0.31
College	0.44	0.42	0.35
Extensive margin	-0.11	-0.07	-0.11
Non-college	-0.13	-0.08	-0.13
College	-0.09	-0.06	-0.09
Intensive margin	0.55	0.48	0.44
Non-college	0.56	0.48	0.44
College	0.53	0.48	0.44

A permanent compensated wage change. We now simulate a permanent-compensated wage decrease of 10% across model economies, which provides an estimate of the Hicks-elasticity of labor supply. Table 10 shows that the Hicks-elasticity of aggregate labor supply does not vary much across model economies. The response along the extensive margin has a different sign from the one along the intensive margin. Most of the variation in the compensated elasticity of aggregate labor supply across model economies is due to the intensive margin. The extensive margin elasticity in the baseline economy and in the re-calibrated economy with $\sigma = 3$ is -0.11. Also, as discussed for the baseline economy, the negative aggregate response along the extensive margin is mostly due to an increase in labor supply of old individuals. Finally, it is also worth pointing that in all economies labor supply responses to a permanent-compensated wage change are much smaller than the ones to a temporary wage change.

The Iceland tax experiment. There are not substantial differences in the aggregate labor supply response when simulating the Icelandic tax experiment in the three economies considered. The re-calibrated economy with $\sigma = 3$ implies an elasticity of 0.57, which is somewhat smaller than the 0.64 value obtained in the baseline economy. The re-calibrated model economy with linear earnings implies an elasticity of 0.54.

7 Conclusions

The inconsistency between the small labor supply elasticities estimated in micro studies and the large ones used in macro models has generated substantial debate among economists. In this paper we develop a quantitative life-cycle theory to study aggregate labor supply responses in a framework that incorporates (i) wage risk, (ii) employment risk and labor market frictions, (iii) heterogeneity in tastes and skills, (iv) tied wage-hours offers, (v) fixed costs of work, (vi) credit constraints and incomplete markets, (vii) government programs (tax and transfer programs). All of these features have been considered as important determinants of labor supply decisions, as discussed in the recent surveys by Keane (2011) and by Keane and Rogerson (2011, 2012). We show that the theory captures a large number of salient features of individual labor supply over the life cycle, by education, both along the intensive and extensive margins. The model is used to study the aggregate labor supply responses to changes in the economic environment. We find that the aggregate labor supply elasticity to a transitory wage shock is 1.75, with the extensive margin accounting for 62% of the response. Furthermore, we find that the aggregate labor supply elasticity to a permanent-compensated wage change is 0.44.

Through a series of experiments aimed at assessing the importance of various modeling assumptions, we find that the intensive margin, although quantitatively important, matters little for understanding the differences in labor supply elasticities across model economies. The impact of modeling assumptions on labor supply responses is almost entirely driven by its effects on decisions along the extensive margin. When we simulate employment responses to a small temporary wage change in an economy with linear earnings, we find that the calibrated fixed costs of work are sufficiently small that, on their own, they do not affect employment decisions along the extensive margin. However, fixed costs of work matter importantly in the presence of non-linear earnings, accounting for half of the employment response in our baseline economy: When we shut down fixed costs of work in the baseline economy, the extensive margin elasticity to a temporary wage change drops from 1.08 to 0.57. Alternatively, we calibrate a version of our economy with linear earnings and find that the extensive margin response to a small temporary wage increase is 2.5 times smaller than in the baseline economy (with linear earnings), even though fixed cost of work are about a a factor of four higher in the former economy. Our results underscore that there is a *complementarity* between fixed costs of work and non-linear earnings that *enhances* the aggregate labor supply response to a temporary wage change.

References

Aaronson, D. and E. French (2004). The effect of part time work on wages: Evidence from social security rules. *Journal of Labor Economics 22*(2), 329–252.

Aaronson, D. and E. French (2009). Deconstructing life cycle expenditures. *Journal of Political Economy 121*(3), 386 – 408.

Aguiar, M. and E. Hurst (2005). Consumption versus expenditure. *Journal of Political Economy 113*(5), 919–948.

Aguiar, M. and E. Hurst (2013). Deconstructing life cycle expenditure. *Journal of Political Economy 121*(3), 437–492.

Bhattacharya, J. and D. Lakdawalla (2006). Does medicare benefit the poor? *Journal of Public Economics 90*(1-2), 277–292.

Bianchi, M., B. R. Gudmunsdsson, and G. Zoega (2001). Iceland's natural experiment in supply-side economics. *American Economic Review 91*(5), 1564–1579.

Blundell, R., A. Bozio, and G. Laroque (2011). Extensive and intensive margins of labour supply: Working hours in the us, uk and france. IFS Working Paper 01/11, Institute for Fiscal Studies.

Chang, Y. and S.-B. Kim (2006). From individual to aggregate labor supply: A quantitative analysis based on a heterogeneous agent macroeconomy. *International Economic Review 47*(1), 1–27.

Chetty, R., A. Guren, D. Manoli, and A. Weber (2011). Does indivisible labor explain the difference between micro and macro elasticities? A meta-analysis of extensive margin elasticities. Working Paper No. 16279, Natonal Bureau of Economic Research.

Cogan, J. (2001). Fixed costs and labor supply. *Econometrica 49*(4), 945?963.

Cooley, T. F. (Ed.) (1995). *Frontiers of Business Cycle Research*. Princeton: Princeton University Press.

Domeij, D. and M. Flodén (2006). The labor-supply elasticity and borrowing constraints: Why estimates are biased. *Review of Economic Dynamics 9*, 242–262.

Domeij, D. and J. Heathcote (2004). On the distributional effects of reducing capital taxes. *International Economic Review 45*(2), 523–554.

Erosa, A., L. Fuster, and G. Kambourov (2012). Labor supply and government programs: A cross-country analysis. *Journal of Monetary Economics 59*(1), 84–107.

Fernández-Villaverde, J. and D. Krueger (2007). Consumption over the life cycle: Facts from the consumer expenditure survey data. *The Review of Economics and Statistics 89*(3), 552–565.

French, E. (2005). The effects of health, wealth, and wages on labor supply and retirement behavior. *Review of Economic Studies 72*(2), 395–427.

Gomme, P., R. Rogerson, P. Rupert, and R. Wright (2004). The business cycle and the life cycle. In M. Gertler and K. Rogoff (Eds.), *NBER Macroeconomics Annual*. The MIT Press.

Gruber, J. and D. A. Wise (Eds.) (1999). *Social Security Programs and Retirement around the World*. Chicago: University of Chicago Press.

Gustman, A. L. and T. L. Steinmeier (1985). The effect of partial retirement on the wage profiles of older workers. *Industrial Relations 24*(2), 257–265.

Gustman, A. L. and T. L. Steinmeier (1986). A structural retirement model. *Econometrica 54*(3), 555–584.

Hansen, G. (1985). Indivisable labor and the business cycle. *Journal of Monetary Economics 16*, 309–327.

Heathcote, J., K. Storesletten, and G. L. Violante (2014). Consumption and labor supply with partial insurance: An analytical framework. *American Economic Review 104*(7), 1–52.

Hornstein, A. and E. C. Prescott (1993). The firm and the plant in general equilibrium theory. In R. Becker, M. Boldrin, R. Jones, and W. Thomson (Eds.), *General Equilibrium, Growth, and Trade II: The Legacy of Lionel McKenzie*. San Diego: Academic Press.

Huggett, M., G. Ventura, and A. Yaron (2011). Sources of lifetime inequality. *American Economic Review 101*(7), 2923–2954.

Imai, S. and M. P. Keane (2004). Intertemporal labor supply and human capital accumulation. *International Economic Review 45*(2), 601–641.

Jaimovich, N. and H. Siu (2009). The young, the old, and the restless: Demographics and business cycle volatility. *American Economic Review 99*(3), 804–826.

Kaplan, G. (2012). Inequality and the lifecycle. *Quantitative Economics 3*(3), 471–525.

Kaplan, G. and G. L. Violante (2010). How much consumption insurance beyond self-insurance? *American Economic Journal: Macroeconomics 2*(4), 53–87.

Keane, M. P. (2011). Labor supply and taxes: A survey. *Journal of Economic Literature 49*(4), 961–1075.

Keane, M. P. and R. Rogerson (2011). Reconciling micro and macro labor supply elasticities: A structural perspective. Working Paper No. 17430, Natonal Bureau of Economic Research.

Keane, M. P. and R. Rogerson (2012). Micro and macro labor supply elasticities: A reassessment of conventional wisdom. *Journal of Economic Literature 50*(2), 464–76.

Keane, M. P. and K. I. Wolpin (2001). The effect of parental transfers and borrowing constraints on educational attainment. *International Economic Review 42*(4), 1051–1103.

Krueger, D. and F. Perri (2006). Does income inequality leads to consumption inequality? evidence and theory. *Review of Economic Studies 73*, 163–193.

Krusell, P., T. Mukoyama, R. Rogerson, and A. Sahin (2011). A three state model of worker flows in general equilibrium. *Journal of Economic Theory 146*(3), 1107–1133.

Low, H. (2005). Self-insurance in a life-cycle model of labour supply and savings. *Review of Economic Dynamics 8*(4), 945–975.

Low, H., C. Meghir, and L. Pistaferri (2010). Wage risk and employment risk over the life cycle. *American Economic Review 100*(4), 1432–1467.

McDaniel, C. (2007). Average tax rates on consumption, investment, labor and capital in the oecd 1950-2003. mimeo, ArizonaStateUniversity.

Meghir, C. and L. Pistaferri (2004). Income variance dynamics and heterogeneity. *Econometrica 72*(1), 1–32.

Meghir, C. and L. Pistaferri (2011). Earnings, consumption and life cycle choices. In D. Card and O. Ashenfelter (Eds.), *Handbook of Labor Economcis, 4B*. Amsterdam: Elsevier.

Moffitt, R. (1984). The estimation of a joint wage-hours labor supply model. *Journal of Labor Economics 2*(4), 550–566.

Osuna, V. and J.-V. Ríos-Rull (2003). Implementing the 35 hour workweek by means of overtime taxation. *Review of Economic Dynamics 6*(1), 179–206.

Pijoan-Mas, J. (2006). Precautionary savings or working longer hours? *Review of Economic Dynamics 9*(2), 326–352.

Prescott, E. (2004). Why do americans work much more than europeans? *Quarterly Review Federal Reserve Bank of Minneapolis*, 2–13.

Rogerson, R. (1988). Indivisible labor, lotteries and equilibrium. *Journal of Monetary Economics 21*, 3–16.

Rogerson, R. and J. Wallenius (2009). Micro and macro elasticities in a life cycle model with taxes. *Journal of Economic Theory 144*, 2277–2292.

APPENDIX

A Hornstein-Prescott Theory of Nonlinear Earnings

We consider first an environment without labor market frictions.

Technology. There are a large number of plants, and each plant is a collection of jobs. We assume that plants can operate jobs at zero costs. The production function of a job is given by

$$f(K,h,z) = h^{\varepsilon} K^{1-\theta} z^{\theta}, \quad \text{with } \theta \leq \varepsilon \leq 1 \tag{A-1}$$

where h denotes the workweek, K is the amount of capital for the job, and z is effective labor in the job (which is given by the worker productivity). Note that, for a fixed workweek, the job technology exhibits constant returns to scale in capital and effective labor. Moreover, as discussed in Osuna and Ríos-Rull (2003), when $\varepsilon = \theta$ the job technology reduces to the standard Cobb-Douglas technology where total hours of effective labor is what matters. When $\varepsilon > \theta$ the hours and effective labor are imperfect substitutes and the composition between these two inputs matters. When $\varepsilon = 1$ the technology is linear in hours and corresponds to the case where workers are not subject to fatigue. While the production function of a job features increasing returns in the three inputs, Hornstein and Prescott (1993) show that the aggregate technology set is convex. Intuitively, the economy's output doubles when the productive resources in the economy double (labor force and capital).

The plant's problem. The plant takes as given the earnings schedule $w(z,h)$ and the interest rate r. For each job, the plant chooses hours of work h, capital K, and effective labor z. In equilibrium, capital is paid its marginal product, and the competition for labor implies that workers will be the residual claimants on the output which remains after capital has been paid and that profits will be zero. Moreover, the earnings schedule is a nonlinear function of the workweek h and a linear function of effective labor z. To show this point, consider a job hiring a worker, with z units of effective labor, for h hours. The optimal amount of capital K solves

$$\pi = \max_{K}\{h^{\varepsilon} K^{1-\theta} z^{\theta} - K(r+\delta) - w(z,h)\}. \tag{A-2}$$

Then, the solution to this problem implies that

$$\frac{K}{z} = k^{*}(r,h) = \left[\frac{(1-\theta)h^{\varepsilon}}{r+\delta}\right]^{1/\theta}. \tag{A-3}$$

Next, notice that a job is open only if profits are non-negative. Free entry, and the fact that jobs can be opened at zero cost, imply that in equilibrium plants will make zero profits from each job. Hence, competition for workers implies that the wage bill $w(z,h)$ is determined from

$$\pi = h^{\varepsilon}\left[zk^{*}(r,h)\right]^{1-\theta} z^{\theta} - zk^{*}(r,h)(r+\delta) - w(z,h) = 0, \tag{A-4}$$

which gives

$$w(z,h) \quad = \quad z(r+\delta)\frac{\theta}{1-\theta}\left[\frac{(1-\theta)}{r+\delta}\right]^{1/\theta} h^{\frac{\varepsilon}{\theta}} = z\Theta h^{\frac{\varepsilon}{\theta}}. \tag{A-5}$$

It follows that the earnings schedule $w(z,h)$ is linear in effective labor z and nonlinear in hours of work h. When $\varepsilon = \theta$ earnings are also linear in h. When $\varepsilon > \theta$ earnings increase more than proportionally with h and, hence, the hourly wage rate also increases with h. In this case, households would be better off by selling employment lotteries to firms (Hornstein and Prescott (1993)). However, we rule out this possibility by assuming that households cannot commit to work when the realization of the employment lottery implies that they should work.

Notice that the log hourly wage of an individual with labor productivity z and working h hours is

$$\ln w_h(z,h) \equiv \ln \frac{w(z,h)}{h} = \ln \Theta + \left(\frac{\varepsilon}{\theta} - 1\right)\ln h. \tag{A-6}$$

This is the functional form typically estimated in the empirical literature on nonlinear wages (see the discussion in Aaronson and French (2009)). The theory presented above provides a theoretical rationale for the functional form used in these empirical studies. Moreover, this is relevant since our calibration strategy will use the estimates from Aaronson and French (2004) to pin down $\frac{\varepsilon}{\theta}$.

Labor market frictions. To introduce labor market frictions into this framework, we follow Krusell et al. (2011) and assume that there are two islands in the economy: a Leisure island and a Work island. Non-employed workers are in the Leisure island and with probability p they move into the Work island. Individuals in the Work island can contact all firms in the Work island at zero cost. Then, competition for workers ensures that payments to workers (in the Work island) are equal to output minus payments to capital. Hence, firm owners make zero profits. This is consistent with equilibrium because jobs in the Work island can be created costlessly. Workers that separate from jobs (exogenously or endogenously) move to the Leisure island.