The (Un)Fortunate Poor

Occupational Sorting with Financial Frictions*

JAN EECKHOUT — UPF Barcelona†

RORY MCGEE — Western Ontario‡

ALIREZA SEPAHSALARI — Bristol§

PETER SPITTAL — Bristol¶

February 15, 2025

— PRELIMINARY AND INCOMPLETE —

— PLEASE DO NOT CIRCULATE OR CITE —

Abstract

We document a robust fact about the working poor: those who are poorer (those with more negative wealth) tend to have higher average wages. Average wages decrease in wealth when wealth is negative, yet they increase for those with positive wealth. This V-shaped pattern with a kink of wages in wealth is robust after controlling for observable characteristics, including age, race, education, and the type of debt — as well as across data sets from many different countries. We propose a theory rooted in occupational sorting with a mean-variance tradeoff: risky jobs pay on average higher wages. We show that in the presence of financial frictions à la Aiyagari, this leads to the V-shaped pattern that we document. In the presence of uninsurable wage risk, in addition to precautionary saving, households also engage in *precautionary occupational sorting*. Those with higher wealth are sufficiently insured against negative outcomes, and they choose riskier jobs. Instead, the poor with negative wealth *gamble for resurrection* hoping to be fortunate, while defaulting on their interest payments if unfortunate. This leads people who have more negative wealth to choose risker occupations and, as a result, receive higher average wages. We quantify the model and evaluate the welfare implications of different redistributive policies.

Keywords. Macro/Labor. Financial Frictions. Precautionary Occupational Sorting. The Working Poor. Wealth Inequality. Wage Inequality.

*We are grateful for comments and feedback from colleagues and seminar audiences. Eeckhout acknowledges support from the ERC, Advanced grant 882499, and from grant PID2022-138443NB-I00 from the Spanish Ministry of Science and Innovation.

†jan.eeckhout@upf.edu – ICREA-BSE-CREI

‡rmcgee4@uwo.ca

§alireza.sepahsalari@bristol.ac.uk

¶peter.spittal@bristol.ac.uk

1 Introduction

Are the working poorest those with the least wealth or those with the lowest wages? We document that those with the lowest wages do not have the least wealth. There is a non-monotonic relation between averages wages and wealth: average wages increase in wealth when wealth is positive (savings), whereas average wages decline in wealth when wealth is negative (debt). In this paper we analyze the relation between wages and wealth of the poor and investigate the underlying economic mechanism. We find that workers sort into occupations with different risk-reward profiles depending on whether they have savings or debt. The inability to insure against wage and wealth shocks leads to precautionary occupational sorting. We quantify the model and find that this has far-reaching welfare implications for commonly used redistributive policies.

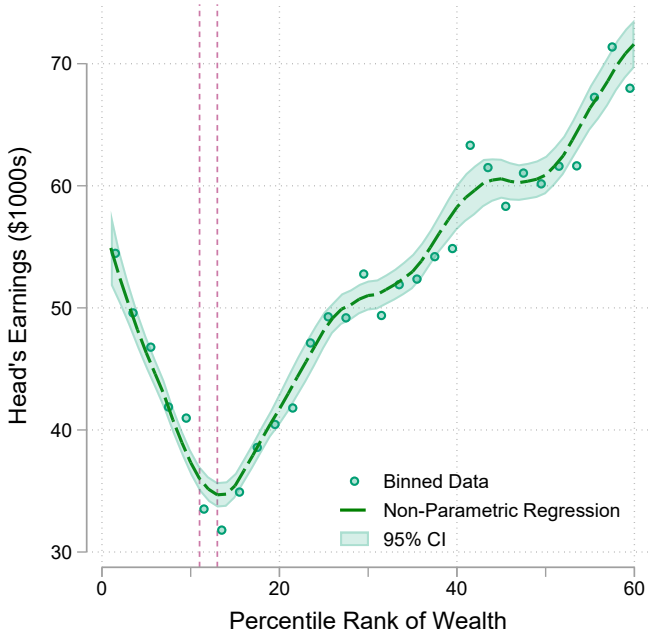


Figure 1: Non-monotone relationship between wealth and annual wages

Notes. This is the estimated kernel regression in equation (1) below, along with 95 percent confidence intervals and a binned scatter plot of the raw data where each point represents average wages in a two-percentile interval on the wealth distribution. The vertical lines delimit the range where wealth is exactly zero.

To investigate the determinants of wages and wealth for the working poor, we document three stylized facts. The first and main fact is that there is a non-monotonic relationship between wealth and wages. Figure 1 plots the relation between household wealth and annual wages. A household head who has no wealth or debt makes on average $35,151. Instead, someone who is in the bottom wealth percentile (with median debt of $169,982) makes $54,909 on average, nearly 60% more, and almost as

much as a household head who is in the 37th wealth percentile ($54,939 on average with median wealth of $53,354).

Our main fact is remarkably robust. The non-monotone relationship is unaltered when we consider different specifications of this relation, including alternative measures of labor income and of wealth. The result is also robust over time and when conditioning on observables and demographics of the households. And we document that the fact also holds for a range of other countries with vastly different labor market and social insurance institutions including Chile, Denmark, Finland, Germany, Norway and the UK.

The second fact establishes that there is also a non-monotonic relationship between wealth and occupations. When we rank occupations by the average wage, then the households with higher positive wealth (more savings) choose higher-paid occupations, whereas those who have higher wealth when wealth is negative (less debt), work in lower-paid occupations. This indicates that there is occupational sorting of households in their wealth holdings.

The third fact shows a positive relationship between the mean and the variance of within-occupation wages. Occupations that pay higher wages also have a higher variance of the wages. This positive correlation between the mean and the variance of within-occupation wages suggests that workers who are risk averse face a mean-variance tradeoff when choosing an occupation.

These facts together help us provide an underlying economic mechanism explaining what happens to the working poor. We propose a theory of occupational sorting in the presence of financial frictions à la Aiyagari (1994). Workers face stochastic labor market outcomes and can choose occupations with different risk-return tradeoffs. In the absence of insurance markets, the wealthy choose high paying occupations because they have enough wealth to self-insure against the low wage outcomes of the high risk, high average wage occupations. Households not only use the canonical *precautionary savings* motive to smooth consumption, they also optimize the occupational choice decision, the *occupational sorting* motive.

The poor with zero wealth choose the occupations with the lowest average wage and the lowest risk. But all poor are not the same. The poor with debt (negative wealth) may end up with such low (or zero in the case of unemployment) wages that they default on the interest they owe on their debt. The threat of default changes the occupational sorting decision. Now the worst wage outcomes, however bad, no longer affect consumption as the household consumes the minimum once it hits the credit constraint. This effectively truncates the wage distribution: it increases the mean and lowers the variance. As a result, the poor with debt gamble for resurrection and prefer those higher ranked occupations once truncated by the possibility of default.

We write down an otherwise canonical Aiyagari (1994) model with two additions. First, there are multiple occupations with different risk-reward tradeoffs and households can occasionally choose a new occupation; Second, households with debt that have a wage realization low enough so they cannot repay their interest, default on their interest payment. The combination of occupational sorting with financial frictions lead to a non-monotonic relation between wages and wealth (fact 1) and between occupations and wealth (fact 2). The positive relation between the mean and variance of wages across occupations (fact 3) is a necessary feature to obtain these findings. We also show that both occupational sorting and financial frictions are necessary conditions to obtain the non-monotonicity. If we shut down either, the main result disappears.

Our objective is to quantify the model matching key moments from the joint wage-wealth distribution using indirect inference, and perform counterfactual experiments on the estimated model, investigating the impact of common policy interventions.

Related literature. Our findings are related to the work by Lewis, Melcangi, and Pilossoph (2019) and Koşar, Melcangi, Pilossoph, and Wiczer (2023) who study the impact of COVID stimulus payments on debt repayment and find that this affects the Marginal Propensity to Consume (MPC).

Daminato and Pistaferri (2023) find a non-monotonic relation of the expected return to net wealth in a household's wealth. That is, households that are in the top percentiles as well as those in the bottom percentiles of the wealth have higher returns on their wealth than the households with returns in the middle of the distribution.

Groes, Kircher, and Manovskii (2015) documents a non-monotonic relation of the job switching probability in the wages. High and low paid workers are more likely to switch job than those paid around the median. The rationale for this pattern is rooted in job search and not in consumption smoothing.

The notion that choices in the labor market offer risk averse workers in financially constrained economies the opportunity to smooth consumption is explored in other work: Eeckhout and Sepahsalari (2024) and Baley, Figueiredo, Mantovani, and Sepahsalari (2022) analyze precautionary job search and mobility decisions, and Huang and Qiu (2022) study precautionary mismatch.

In the theory literature, Stiglitz and Weiss (1981); Sinn (1982); Gollier, Koehl, and Rochet (1997) analyze the role of limited liability constraints for firm investment. Gollier, Koehl, and Rochet (1997) find that firm size is too large, and when they have the choice of investment opportunities with different mean-variance assets, firms may choose assets with mean-variance characteristics that are inefficient. Sinn (1982) analyzes shocks to human wealth, such as unemployment and health, have a lower bound

which can be zero or positive when society has arrangements to protect poverty. The paper argues that this leads to limited liability and hence the inefficient choice of insurance policies. And the work by Stiglitz and Weiss (1981) studies the role of equilibrium interest rates in sorting loan applicants: those who pay higher rates tend to pick projects that have a higher payoff but also higher risk of default.

2 Data and Empirical results

2.1 Data

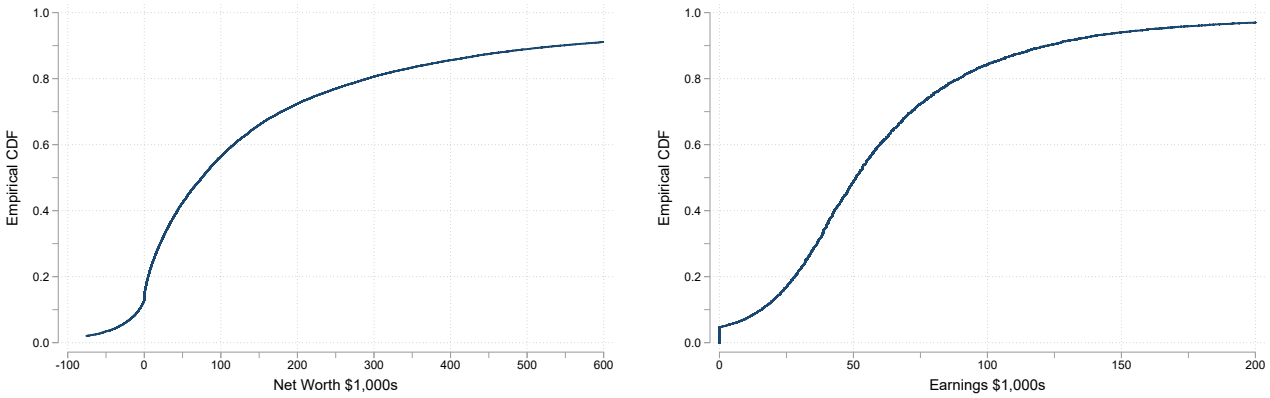
Our core data are drawn from the Panel Study of Income Dynamics (PSID), a biennial household panel survey containing detailed information on wages, wealth and demographic characteristics for a nationally-representative sample of households in the US. We focus on the period from 1999 to 2019 because, corresponding to the period over which PSID has collected asset data in every wave.

We also document our key facts for a set of other countries where high-quality data on wealth and wages are available to us. These additional data are drawn from two sources. For the UK, we use the UK Household Longitudinal Study (UKHLS), which contains high-quality data on wealth, wages and background characteristics for individuals from a representative sample of households; for the other countries we study—Austria, Australia, Chile, Denmark, Estonia, Finland, Greece, Italy, Luxembourg, Germany, Norway, Slovenia, Slovakia, South Korea, Spain and Sweden—we use data from the Luxembourg Wealth Study. Further details on these additional data and how we have processed them are available in Appendix A.

Sample selection. To construct our baseline sample, we begin with an unbalanced panel of all PSID individuals in the sample period. We then focus only on married households, and restrict further to those with a head aged between 25 and 55 to abstract from education or retirement decisions. Following Blundell, Pistaferri, and Saporta-Eksten (2016), whenever there is a change in marital composition we drop the household in the year of the change and treat it as a new family in subsequent observations. Finally, to reduce the influence of measurement error, we drop 33 observations with extremely high asset values ($20 million or more). Our final sample contains a total of 26,383 employed observations with non-zero survey weights, comprised of 6,552 households observed for an average of 4.03 waves.

Wealth. Our primary measure of wealth is intended to capture households' net worth. It is defined as the household's total assets — that is, the sum of cash, bonds, stocks, the value of any business, pension

funds, housing, and other real estate, and vehicles — net of any mortgage or other debt.[1] We adjust our measure of wealth to 2012 prices using the CPI to adjust for inflation.



(a) Wealth (b) Wages

Figure 2: Empirical Distributions of Wealth and Wages

We report the empirical CDF of wealth in Figure 2a. This shows that 12.5 percent of households have negative wealth, and a small mass of 2 percent of households report precisely zero wealth. Above this, the distribution has a long tail: half of households report wealth above around $90,000, while around 10% report assets in excess of $600,000. To reduce the influence of households with extremely high or low wealths on our results, we use the percentile rank of wealth in our data (rather than its level) throughout our analysis.[2]

Wages. Our measure of wages is the household head's gross annual labor income. This captures wages from all jobs, plus any labor component of income from a business or farm, inflated to 2012 prices using CPI. Wages are thus the annual wage rate, and not the hourly or weekly wage rate. We also use data on the occupation of the household head's primary job using census occupation codes provided by the PSID.[3] We report the empirical CDF of wages in Figure 2b. Following, Guvenen, Karahan, Ozkan, and Song (2021), who study income dynamics in US administrative social security

[1] We also consider a more restrictive measure of wealth, "liquid" wealth, which is intended to capture resources that are easier for households to access or borrow against on short notice. Liquid wealth excludes the value of home equity (i.e., the net value of a primary residence after subtracting any mortgage), the value of vehicles, and the value of pension funds. All our main findings are robust to the choice of asset measure.

[2] Constructing percentile ranks separately by PSID wave allows us to flexibly control for changes to the distribution of wealth over time. But it means that the wealth value corresponding to any percentile rank is not fixed over time. For example, a household in wave t with a percentile rank of 50 was at the median of the wave-t wealth distribution, but the wealth holdings corresponding to this may be different from a household at rank 50 in wave $t+1$. Results are robust to this alternative definition.

[3] The census determines occupations, and the granularity of occupation data varies across waves of the PSID. However, when we construct occupation-specific, outcomes we average over, or rank, occupations in each different occupation code regime. This is a non-parametric approach that accounts for variation in the occupation codes as well as occupation specific secular trends.

records, we treat reported wages less than one quarter of full-time work (40 hours per week) at half of the legal federal minimum wage in year t as zero.

2.2 Results

In this section, we document three key facts on the relationship between wealth and labour market outcomes. First, we show that there is a clear and robust non-monotonicity in the relationship between household wealth and average wages: wages initially fall as wealth increases, before reaching an inflection point and starting to increase. This fact holds for the US and a broad of other countries – with very different institutional arrangements – for which we have access to high-quality data on wealth and wages.

Second, we document a similar non-monotonic relationship in the relationship between wealth and occupation choice. That is, the average pay in a primary earner's occupation follows a very similar relationship with wealth as their own average wages. This suggests that at least part of the non-monotonicity in wages is driven by selection into occupations with different average pay

Finally, we document that there is a positive correlation across occupations between the mean and variance of pay. Given that average occupational pay follows a non-monotonic relationship with wealth, this implies that wage dispersion should also follow a similar pattern – and we show that, indeed, it does.

Fact 1: Non-monotonic relationship between wealth and wages

We begin by estimating the following non-parametric relationship between wages, $y_{i,t}$, and the household's rank in the time-t asset distribution $a_{i,t}$:

$$y_{i,t} = m(a_{i,t}) + \epsilon_{i,t}, \tag{1}$$

where the unknown function $m(\cdot)$ defines the mean of wages conditional on wealth. We estimate this function using local linear regression with a bandwidth selected by cross-validation.[4]

We present the estimated function in Figure 1, along with 95 percent confidence intervals and a binned scatter plot of the raw data—each point represents average wages in a two-percentile interval on the wealth distribution. The non-parametric function is estimated precisely, as indicated by the tight confidence bands.

[4]We use the same bandwidth across all regressions to aid comparability and to avoid numerical issues in estimating higher order moments when the bandwidth differs. The bandwidth selected by cross validation is 2.31.

Our estimates reveal a surprising pattern: average wages are non-monotonic in wealth. At the bottom of the asset distribution, average wages decline in wealth. But this relationship reverses at around the 15th percentile of the asset distribution—that is, at around the same point in the distribution when wealth equals zero. Beyond this point, average wages increase in wealth.

In Table 1, we demonstrate that the non-monotonic relationship is robust to controlling for observable differences across households. We control for these features within a piecewise-linear regression of the form:

$$y_{i,t} = \beta_0 + \beta_1 r_{i,t} + \delta_0 \mathbf{1}_{a_{i,t}>0} + \delta_1 r_{i,t} \mathbf{1}_{a_{i,t}>0} + \gamma X_{i,t} + \epsilon_{i,t}, \tag{2}$$

where $r_{i,t}$ is the percentile rank of wealth and the non-monotonic relationship between wealth and pay would manifest as a difference between the slopes for positive wealth, δ_1, and negative wealth, β_1. In this specification we additionally control for household-level characteristics $X_{i,t}$.

Table 1 reports our estimates of the parameters of interest. Column (1) reiterates the pattern shown in Figure 1: there is a non-monotonic relationship between wages and wealth, with a statistically significant difference in the relationship between those with positive versus negative wealth. In column (2) we add controls for age and race, and in column (3) we control for education. While these controls attenuate the slope coefficients, the clear non-monotonicity remains statistically significant across all specifications. In the most conservative specification an increase of 1 percentile rank is associated with an almost $1300 increase in wages for those with positive wealth. In contrast, for those with negative wealth a *decrease* of 1 percentile rank is associated with a similar *increase* in wages.

To better understand the drivers of this non-monotonic relationship, we show the relationship between wealth and labour supply on the intensive and extensive margin in Figure 3. First, we note that that there is no visible non-monotonicity between wealth and annual hours of work (conditional on working). Second, Figure 3 shows a statistically significant non-monotonicity in labor supply at the extensive margin, with the probability that the household head is in employment first falling and then increasing again in wealth. However, we note that this kink is quantitatively small, with a difference of around 5% in the probability of employment between the bottom of the wealth distribution and the point where wealth equals zero. This is considerably smaller than non-monotonicity in annual wages conditional on working present in Figure 1, in which wages fall by over 35%.

Taken together, the lack of a clear non-monotonicity in hours of work and the quantitatively small non-monotonicity at the extensive margin suggest that the non-monotonic relationship we document comes from variation in *wages*, not labor supply. This feature of the empirical results helps motivate the model we develop in Section 3.

Table 1: Parametric test for non-monotonicity

	(1)	(2)	(3)
Wealth ($a_{i,t} \geq 0$)	1463.328***	1429.645***	1281.147***
	(123.899)	(123.845)	(121.330)
Wealth ($a_{i,t} < 0$)	-1570.578***	-1462.363 ***	-981.7983***
	(345.543)	(344.331)	(333.278)
Controls			
Age		✓	✓
Race		✓	✓
Education			✓
N	26,383	26,317	26,317

Notes. In all regressions we additionally control for state and year fixed effects. Standard errors clustered at the household level. * $p < 0.10$, ** $p < 0.5$, *** $p < 0.01$

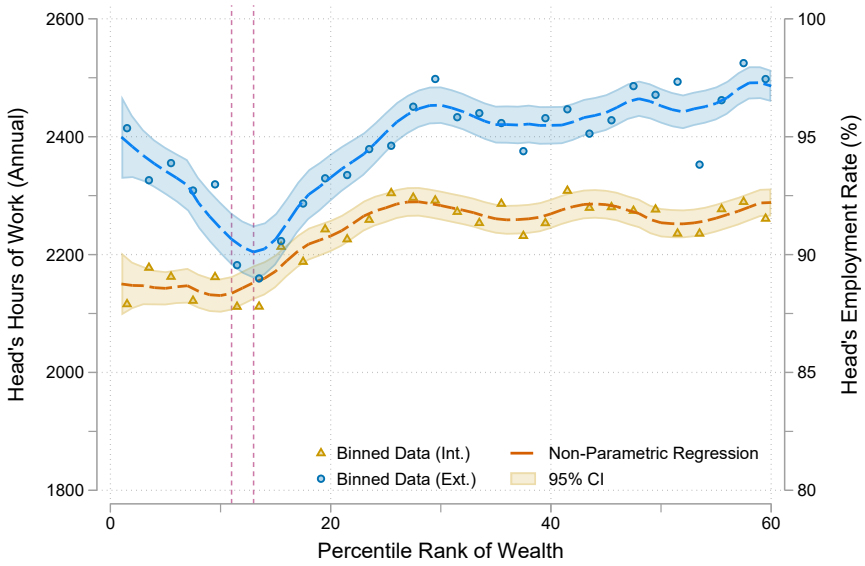


Figure 3: Relationship between wealth and labor supply

Notes. This is the estimated kernel regression in equation (1), with (1) a binary measure for whether worker i is in work at time t, estimated on a sample of all married men aged 25-55, and (2) the annual hours of work among those who are working as dependent variables. For each estimated relationship, we show also 95 percent confidence intervals and a binned scatter plot of the raw data where each point represents average wages in a two-percentile interval on the wealth distribution. The vertical lines delimit the range where wealth is exactly zero.

In Appendix B, we document a number of additional results which underline the robustness of the non-monotonic relationship between earnings and wealth. First, in Appendix B.1, we show a qualitatively similar relationship using a measure of *liquid* wealth instead of total wealth, i.e. the sum of

cash, bonds, stocks, the value of any business, net of any non-housing, non-vehicle debt; or, to put it differently, total wealth minus the value of pension funds, home equity and vehicles.

In Appendix B.2, we show that the pattern is robust to considering alternative labor market outcomes including the percentile rank of annual wages, which is less influenced by individuals with extreme wages; post-tax annual wages, which may better reflect consumption opportunities; a measure of permanent income, taken as the residual from a Mincerian wage regression; and hourly wages.

In Appendix B.3, we demonstrate that our main results are robust to alternative sample definitions, such as considering single people or including those who do not work in our sample. We also document quantitatively similar non-monotonic relationship for each year in our sample separately. Finally, to address concerns that the kink may be driven by the earnings of those who report precisely zero wealth in our data, we show that the non-monotonic relationship persists even when we exclude individuals with wealth in the region of zero.

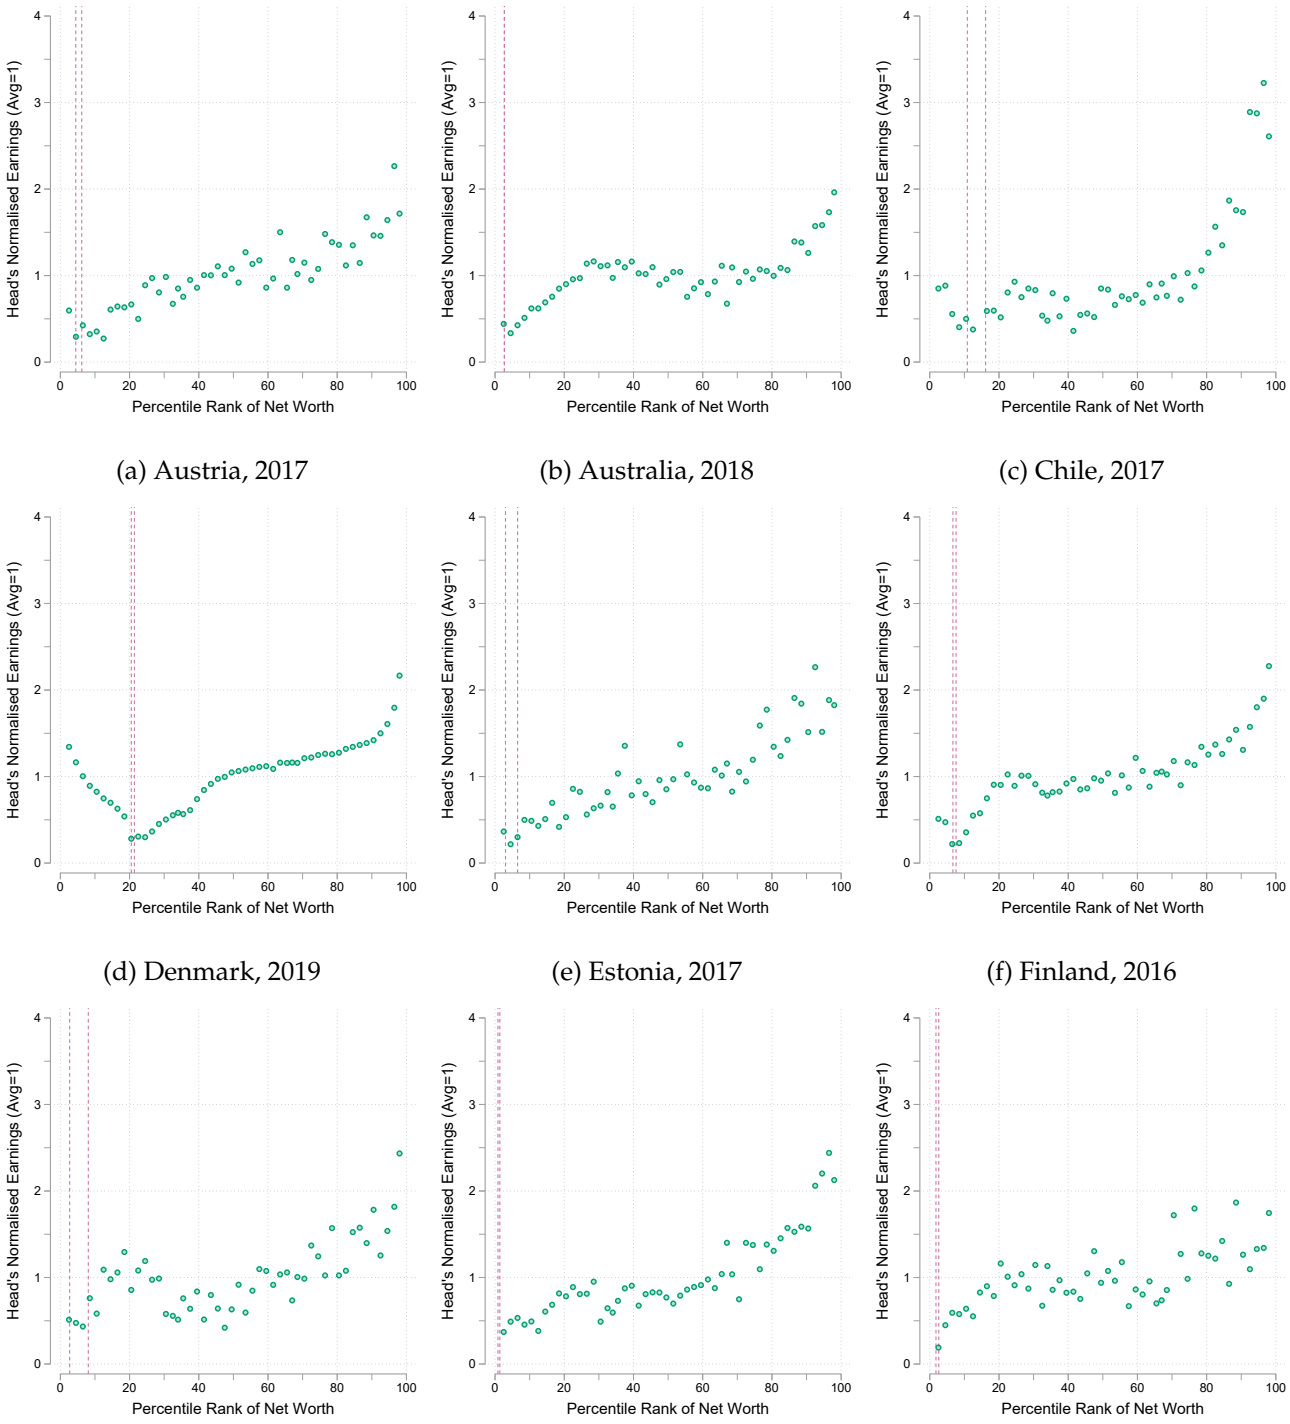
Strikingly, this non-monotonic relationship between wages and wealth is also present across a range of other countries with vastly different labor market and social insurance institutions. Figure 4 and Figure 5 contains bin-scatter plots of average wages (expressed as fractions of the average wage in each country) against the percentile rank of wealth for Austria, Australia, Chile, Denmark, Estonia, Finland, Germany, Greece, Luxembourg, Norway, Slovakia, Slovenia, South Korea, Spain, Sweden and the UK. These are all the countries for which high-quality data on wealth and wages are available to us in either the Luxembourg Wealth Study or the UK Household Longitudinal Study.[5]

As in Figure 1, each point represents average wages in a two-percentile interval on the asset distribution. The vertical lines in each panel indicate the size and position of any mass with zero wealth measured for each country: the left-hand vertical line indicates the percentile rank of an individual with wealth of -100 local currency units, while the right-hand line indicates the position of the individual with +100 local currency units.

For 13 out of 16 of these countries, average wages fall in wealth for those with the lowest wealth—indeed, across all countries, this is for individuals with negative wealth. But this relationship is reversed for those with higher wealth—meaning, in all cases, those with positive wealth. The exceptions are Italy, Luxembourg and Slovenia, for which the underlying data used by the Luxembourg Wealth Study contains few individuals with negative wealth—and so there is little density on the downward-sloping segment of the kink. For countries with the highest quality wealth data, in particular the Scandinavian countries, we see a very clear pattern that is qualitatively identical to the results we document in the

[5]The LWS contains harmonised data on income and wealth, thus we rely on this to draw our comparator group. We use the most recent available dataset with wealth and income available jointly that also overlaps with the sample period for our US data. In addition, to avoid artificially compressing the income distribution in all countries due to a long right tail in a single country, we truncate earnings at four times the national average.

US.

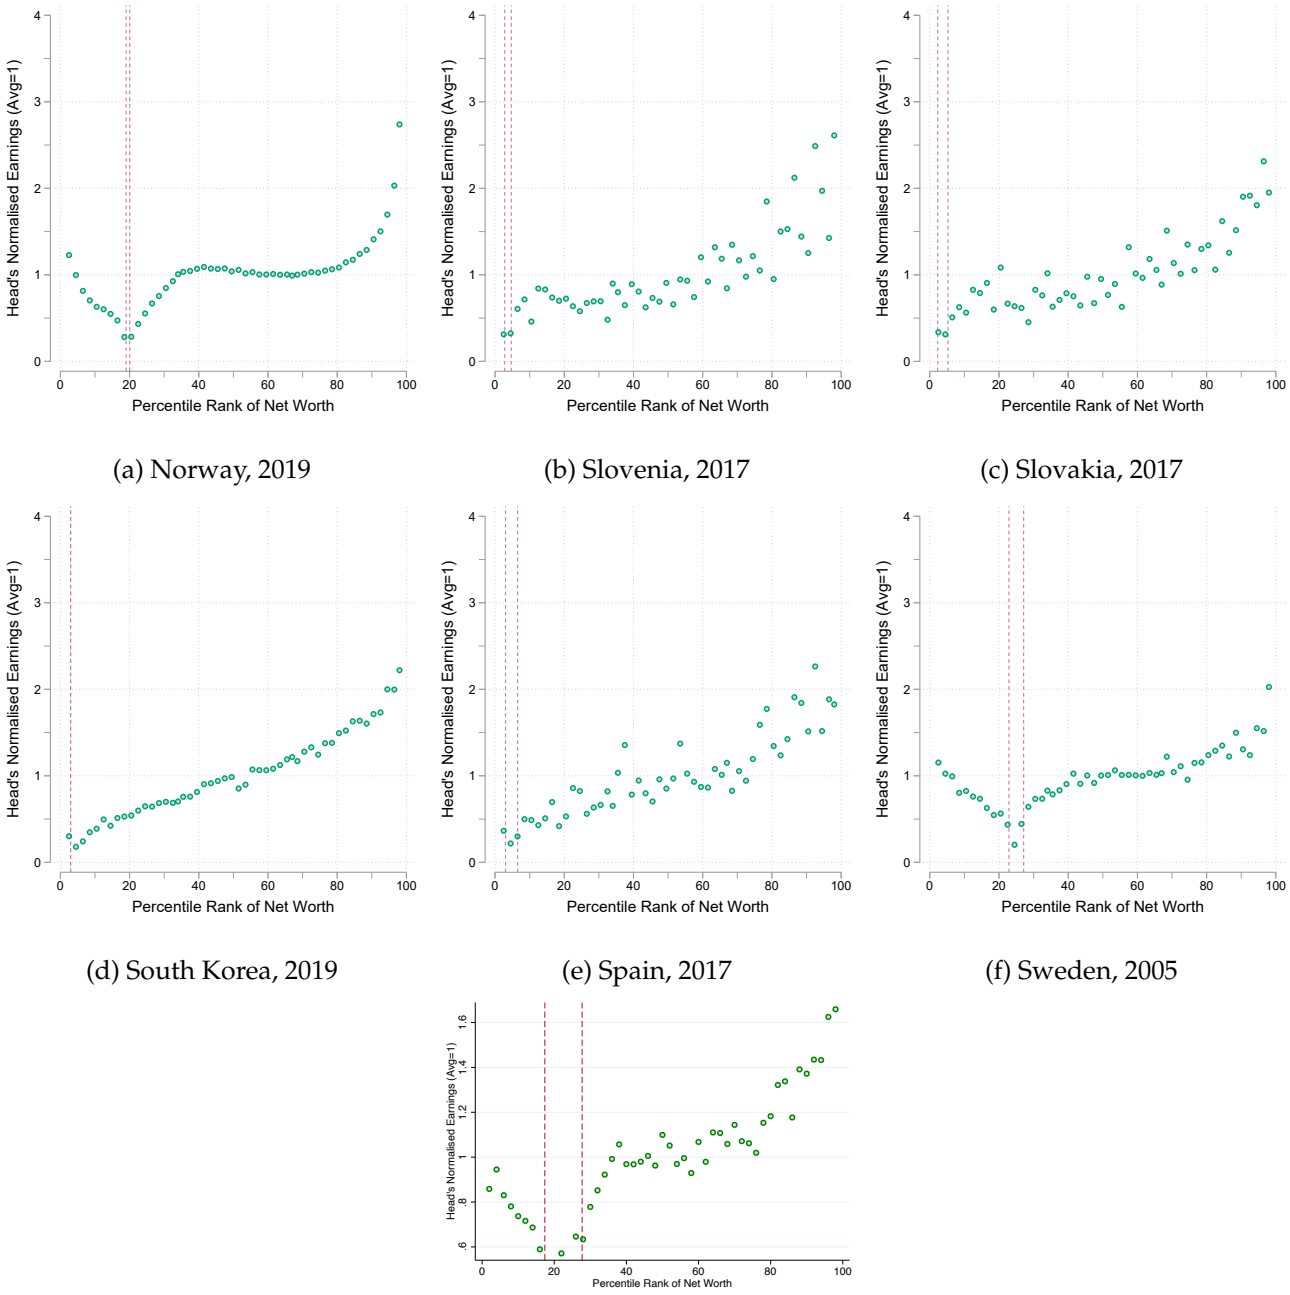


(a) Austria, 2017 (b) Australia, 2018 (c) Chile, 2017

(d) Denmark, 2019 (e) Estonia, 2017 (f) Finland, 2016

(g) Greece, 2018 (h) Italy, 2016 (i) Luxembourg, 2018

Figure 4: Relationship between wealth and wages across countries (I)



(a) Norway, 2019 (b) Slovenia, 2017 (c) Slovakia, 2017

(d) South Korea, 2019 (e) Spain, 2017 (f) Sweden, 2005

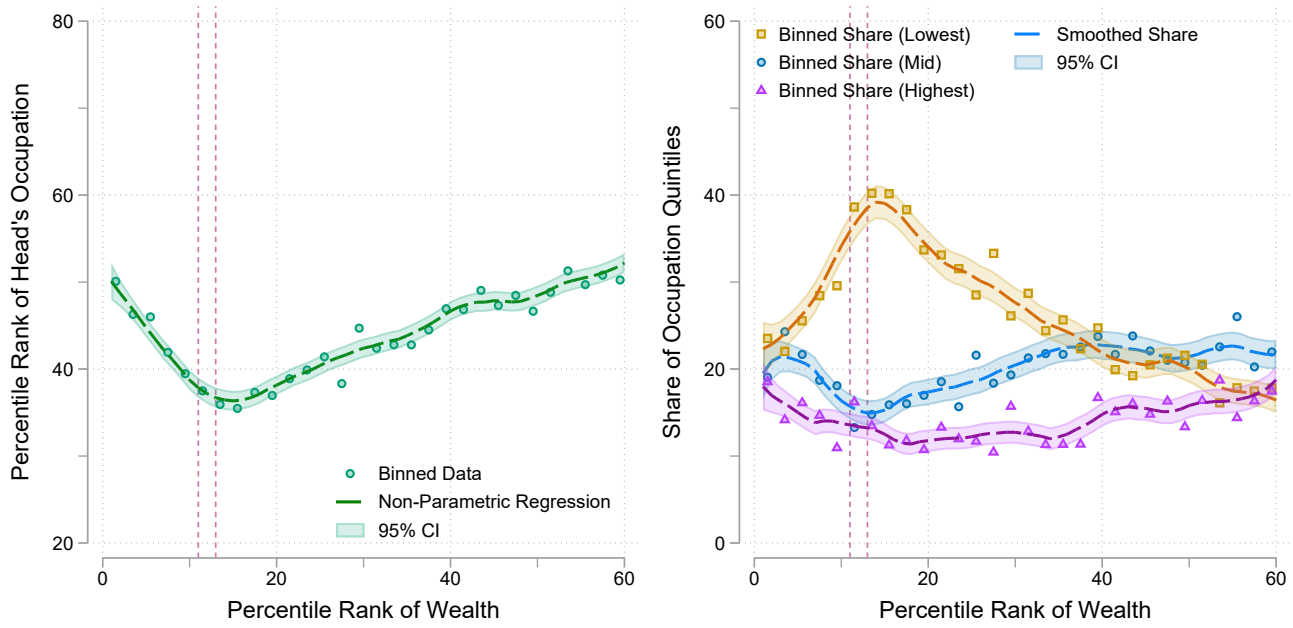
(g) UK, 2017

Figure 5: Relationship between wealth and wages across countries (II)

Fact 2: Non-monotonic relationship between wealth and occupations

We now turn to the role of occupational choice in driving the non-monotonic relationship between wages and wealth. We begin by constructing average wages within each occupation—measured in our data as the census occupation code of the household head's primary job—and then rank occupations from lowest to highest according to their average wage. We rank occupations separately by occupation

code regime to flexibly control for aggregate wage trends within occupations.



(a) Average occupation rank

(b) Distribution across occupations

Figure 6: Non-monotone relationship between wealth and occupation

First, we estimate a similar non-parametric regression as in equation (1), except we redefine $y_{i,t}$ to be the occupation pay rank of household head i at time t. We plot the estimated relationship, which is very similar to that between wages and wealth, in Figure 6a. The heads of households with the very lowest wealth work in occupations which are, on average, at the 51st percentile of the occupation average pay distribution. The occupation rank then initially decreases in wealth, with workers at around the 15th percentile of the asset distribution working in occupations which are, on average, at the 35th percentile of average pay. Beyond this, the average pay of the household head's occupation increases in wealth. This indicates that workers at different parts of the asset distribution do not just earn different amounts, but they are drawn from systematically different occupations.

This point is underlined by Figure 6b. We group occupations into quintiles based on their average pay, and show the probability that a household head's occupation is in the bottom, middle and top quintile as a function of their wealth (we omit the second and fourth quintiles from the figure for legibility). The probability of being in one of the top or middle-paying occupations follow the familiar non-monotonic pattern: first falling in wealth, and then increasing again beyond the 15th percentile of the asset distribution—albeit with a more pronounced kink for the middle than for the top occupations. By contrast, the probability of being in one of the lowest-paying 20 percent of occupations follows the inverse pattern: first rising in wealth, and then falling.

Fact 3: Positive relationship between mean and variance of occupation wages

What drives this striking pattern of occupational sorting? Why do workers initially sort into lower-paid occupations – and as a result receive lower average wages – as their wealth increases? And what leads this relationship to reverse for sufficiently high wealth?

To attempt to answer these questions, we document one further fact: there is a positive relationship between average wages and the dispersion of pay within occupations. We demonstrate this for the US in Figure 7, which contains a scatter plot of the mean and standard deviations of earnings across occupations. There is a clear positive relationship, with a correlation coefficient of 0.795. This indicates that indeed occupations with the highest average pay also have the highest wage dispersion—and are therefore more risky prospects.



Figure 7: Positive relationship between mean and standard deviation of wages across occupations

An implication of this positive relationship between the mean and variance of wages across occupations, combined with the fact that there is a non-monotonic relationship between occupation choice and wealth, is that there should also be a non-monotonic relationship between the variance of pay and wealth.

To establish whether this is the case, we first construct the squared residual $(y_{i,t} - \hat{m}(a_{i,t}))^2$ from regression (1). We then estimate the following non-parametric relationship between wealth and the variance of wages:

$$(y_{i,t} - \hat{m}(a_{i,t}))^2 = v(a_{i,t}) + u_{i,t}. \tag{3}$$

where the function $v(\cdot)$ defines the variance of wages conditional on wealth. We present the estimate function in Figure 8, along with 95 percent confidence intervals and a binned scatter plot of the underlying data.

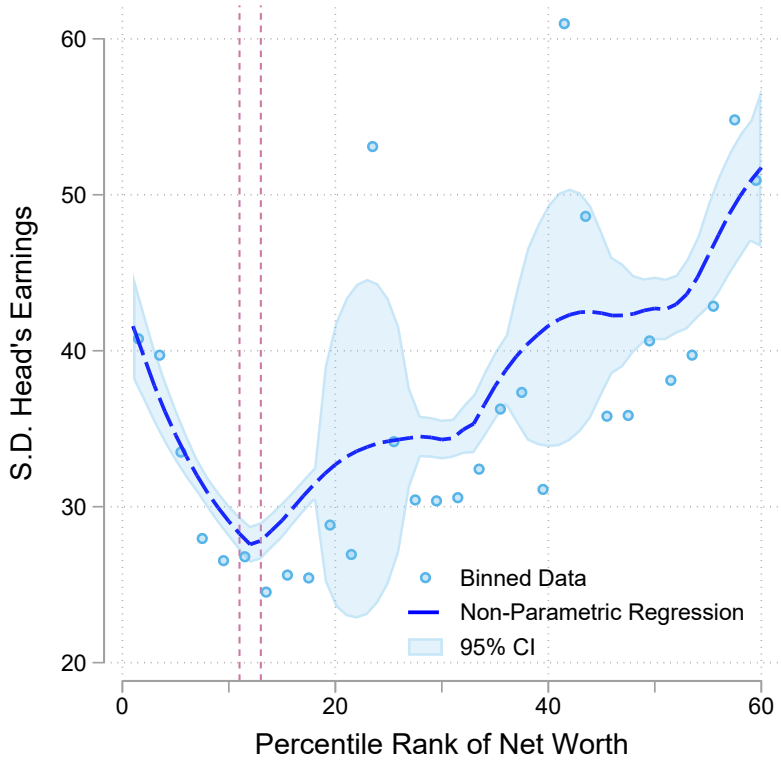


Figure 8: Non-monotonic relationship between wage dispersion and wealth.

Notes. we construct our estimate of the conditional variance using residuals from the relationship estimated in Figure 1. The asymptotic distribution of this estimator is derived in Fan and Yao (1998) and we use the delta method to construct the asymptotic distribution for the conditional standard deviation.

The non-monotonic relationship for the standard deviation of wages in Figure 8 echos the same relationship for the average labor wages in Figure 1. Thus, variances are also non-monotonic in wealth. At the bottom of the asset distribution, the cross-sectional variance of wages decline in wealth. But, at zero, this relationship reverses. In general, the non-parametric function is estimated precisely, as indicated by the tight confidence bands for much of the distribution. While the relationship is less precisely estimated at parts of the distribution with relatively large outliers, the point estimates evidence the non-monotonic pattern throughout. Moreover, we have a high degree of statistical precision at the inflection point and are, thus, confident in the non-monotonic pattern we document.

3 The Model

We build a simple model that captures the risk-return trade-off as a function of wealth through occupational choice. The model accommodates the three facts that we document above: 1. wages are non-monotonic in wealth; 2. occupational sorting is non-monotonic in wealth; 3. there is a monotonic relation between the mean and variance of occupation wages. We do this in a dynamic model where households make a consumption-savings decision in conjunction with occupational choice.

Setup. Time is discrete and indexed by t, and the horizon is infinite. The small, open economy is populated with a continuum of households of measure one. In each period t, households work to earn a wage w and consume. Households face uninsurable wage risk and can borrow and save to smooth consumption. Households are infinitely lived.

Preferences. Denote consumption in period t by c_t, and wealth by a_t. Households are risk averse with preferences represented by a Von Neumann-Morgenstern utility function $u(c)$ over consumption c. We assume that u is twice differentiable, increasing and concave: $u' > 0, u'' < 0$. Agents discount utility by a factor $\beta < 1$. Households can trade a risk-free bond at interest rate $r(a)$ that depends on whether wealth is positive or negative. When assets are positive, they save rate $r(a) = r$. When wealth is negative, they borrow, they borrow at rate $r(a) = r + \delta$, where δ reflects the spread due to default. The interest rate is determined exogenously as in a small, open economy, and so is the spread. Household borrowing is constrained at a threshold, $\underline{a}$. Households may not default on the principal of their debt but, as explained below, may default on interest payments.

Technology. In each period, households supply one unit of labor at wage w. Their total income is this determined by wages and the return on the bond. They are exposed to three sources of idiosyncratic risk: 1. an occupation shock, 2. a wage shock, and 3. an expense shock.

There are n occupations $y \in \{1, 2, ..., n\}$. At the start of each period, workers are employed in an occupation y. The first shock allows households to choose a new occupation. With probability λ, the household can choose occupation, optimizing over the expected stream of wages associated with that occupation. With complementary probability $1 - \lambda$, workers remain in the same occupation $y_t = y_{t-a}$. When $\lambda = 1$ they choose an occupation every period.

The second is the wage shock. The wage $w(y)$ is stochastic and the stochastic process depends on y with $\mathbb{E}w(y)$ and Var $w(y)$ increasing in y: higher y jobs generate a higher expected income, and are

more risky. Wages in each occupation follow an AR(1) process

$$\log\ w_t(y) = \rho_y \log\ w_{t-1}(y) + \varepsilon_t(y), \quad \varepsilon_t(y) \sim \mathcal{N}\left(\mu(y), \sigma(y)\right) \quad \text{i.i.d} \tag{4}$$

where ρ_y is the persistence of the process, and μ_y is mean and σ_y is the standard deviation of innovation. Given the AR(1) property, log wages follow a normal distribution $F(w|y)$ with mean $\frac{\mu_y}{1-\rho_y}$ and variance $\frac{\sigma_y^2}{1-\rho_y^2}$. Reflecting empirical Fact 3, we impose that both are increasing in y. Higher y jobs create higher expected income and are more risky.

The third is an expense shock, which is an idiosyncratic shock to their wealth. In principle this could be part of a more general income process, but we distinguish it from the wage process which we identify from the labor income process. The expense shock is a shock that is independent of the occupation and without persistence. Denote the wealth holdings of a household in period t by a_t. Every period, with probability ϕ workers face a shock to their wealth in which case they lose a part of their wealth $\tilde{a}$. Following Chatterjee, Corbae, Nakajima, and Ríos-Rull (2007), we motivate this feature by assuming that households experience a necessary expense such as medical expenses or an accident.

Default. In this environment, there are two forms of default. First, we assume that if the interest payment $r(a)a_t$ that is due, exceeds the wage w_t, the household defaults on the interest payment. To keep the model tractable, this is not a choice. Default has a cost, and we consider three alternatives: 1. a wage cost where the households can keep a share $\alpha \in [0,1]$ of the wage; a fixed cost κ where the household pays a penalty κ; 3. a utility cost μ. We investigate quantitatively the role of each. Throughout, households that default on their interest payments can continue to save and borrow the risk-free bond. Second, if the expense shock drives the household below the borrowing constraint $\underline{a}$, it defaults on the part the debt in excess of the constraint and ends up with wealth $\underline{a}$.

We assume that there is no insurance market for these idiosyncratic occupation, wage and wealth shocks. In the presence of the borrowing constraint, households cannot fully self-insure since they cannot assume large negative wealth outcomes which occur with positive probability. They can only use the actions available to them and that offer precautionary insurance. They can save or borrow risk-free bonds to insure against these shocks—*precautionary saving*—or they can sort into the different occupations that offer different risk profiles—*precautionary sorting*.

Actions and value functions. The timing of events in any period t is as follows. (1) The occupation shock realizes with probability λ, and the shocked household makes the occupational choice decision y_t. With complementary probability $1-\lambda$ the household makes no occupational choice decision and

$y_t = y_{t-1}$; (2) The wage shock $w_t(y)$ realizes; (3) the wealth shock realizes with probability ϕ, resulting in a loss of wealth $\tilde{a}$; (4) Given past savings, the wage and the wealth shock, households make their consumption-saving decision and consume. They then proceed to the next period. In what follows, we focus on steady-state equilibrium.

The recursive formulation of a household's decision problem can be described by 4 value functions V, W, J:

(1) **Occupational Sorting.** At the start of the period t, the expected value of a household who was working in occupation y_{t-1} last period and has saved wealth a_t is denoted by $V(y_{t-1}, a_t)$:

$$V(y_{t-1}, a_t) = \lambda \max_{y_t} W(y_t, a_t) + (1-\lambda) W(y_t, a_t)|_{y_t = y_{t-1}}. \tag{5}$$

With probability λ the household optimally chooses a new occupation y_t; with complementary probability the occupation remains unchanged and $y_t = y_{t-1}$.

(2) **Wage and Expense shocks.** The wage shock realizes, drawn from the distribution $F(w_t|y_t)$ that is conditional on the current occupation, resulting in the ex post wage w_t. In addition, there is an expense shock: with probability ϕ, assets are reduced by $\tilde{a}$. We can write the ex post assets after the wage and expense shocks denote by a'_t as follows:

$$a'_t(\mathcal{A}, \mathcal{D}) = \max\{a_t + (1-\mathcal{D})((r+\delta)a_t + w) + \mathcal{D}(\alpha w - \kappa) - \mathcal{A}\tilde{a}\ ,\ \underline{a}\} \tag{6}$$

where the indicator $\mathcal{A} \in \{0, 1\}$ is equal to 1 if an expense shock occurs and equal to zero otherwise, and where the indicator $\mathcal{D}(a, w) \in \{0, 1\}$ is equal to one if the household defaults on the interest payment. Interest payment default occurs when $(r+\delta)a_t + w_t < 0$, in which case assets are $\alpha w - \kappa$. In addition, assets never drop below the borrowing constraint $\underline{a}$. We write the continuation value W as:

$$W(y_t, a_t) = \phi \int_{w_t} J(w_t, y_t, a'_t(\mathcal{A}, \mathcal{D}))dF(w_t|y_t) + (1-\phi) \int_{w_t} J(w_t, y_t, a'_t(\mathcal{A}, \mathcal{D}))dF(w_t|y_t). \tag{7}$$

(3) **Consumption-savings decision.** Given the post-shock realization of wealth a'_t, the household chooses consumption c_t and savings/borrowing a_{t+1} with continuation value $V(y_t, a_{t+1})$ and sub-

ject to the budget constraint:

$$J(w_t, y_t, a'_t) = \max_{c_t, a_{t+1}} \{u(c_t) - \mathcal{D} \times \nu + \beta V(y_t, a_{t+1})\} \tag{8}$$
$$\text{s.t. } c_t + a_{t+1} \leq a'_t$$
$$a_{t+1} \geq \underline{a}.$$

Note that a'_t includes the wage w_t as well as the expense shock and default (from equation (6). Observe also that we allow for disutility ν in case the household defaults (in addition to the proportional wage retention $(1-\alpha)w_t$ and the monetary penalty κ).

Equilibrium. With an exogenous interest rate or absent strategic interaction or externalities, an equilibrium is a dynamic decision problem that specifies a sequence $\{y_t, c_t, a_{t+1}\}$ for each household that satisfies the value functions in equations (5), (6), (7), (8), and which results in a steady state distribution of occupations, wages and wealth.

Limit Cases. We simulate the model for two limit economies: 1. no financial frictions; 2. no occupational sorting. We illustrate that in both cases there is no non-monotonicity of expected wages in wealth, which establishes that financial frictions and occupational sorting are necessary in this framework to generate the empirically observed patterns.

1. No financial frictions ($\underline{a} \to -\infty$). Households can perfectly insure against any wage and wealth shocks by holding infinitely negative wealth. As a result, their consumption is constant in every period and equal to the expected wage minus the expected wealth shock. Households choose the occupation $y = n$ with the highest expected wage as they derive no disutility from risky outcomes. In the absence of financial frictions, there is no occupational sorting – more precisely, there is degenerate sorting into the top occupation – resulting in constant expected wages.

2. No occupational sorting ($n = 1$; $y \in \{1\}$). When there is no occupational choice, households with lower wealth and closer to the borrowing constraint will save more and consume less. Any positive shock results in more consumption and hence a lower marginal propensity to consume. Since they cannot switch occupations, they cannot smooth consumption with precautionary occupational sorting.

4 Simulations

We simulate the economy to generate the key properties of the model. As in the first step of a calibration exercise (see Gourinchas and Parker (2002)) where we estimate the parameters of the model we choose similar commonly used parameters. Unlike the calibration exercise, we also externally set step two parameters.

Externally set parameters. We set three parameters (the discount rate, the interest rate and the occupational switching rate, see Table 2) externally, either because there is a direct counterpart in the data, or those parameter values have been widely used in other studies. The model and data period is one year. The discount factor β is chosen to be 0.965 which corresponds to 3.5% annual discounting (see Krusell, Lou, and Ríos-Rull (2019)). We set the annual risk-free rate equal to 2.5%, equal to the average US 10 year treasury yield between 2010 and the pandemic. To calibrate the occupational switching rate, we use the separation rate from the Current Population Survey (CPS) (see Eeckhout and Lindenlaub (2019)).

Table 2: Externally set parameters

	Definition	Value	Source
β	Discount rate	0.965	Krusell, Lou, and Ríos-Rull (2019)
r	Interest rate	2.5%	US Treasury Yield
λ	Occupational switching rate	0.24	Current Population Survey (CPS)
γ	Risk aversion	2.3	
ϕ	Expense shock	0.07	
α	Wage share post-default	1	

Functional forms and other parameters. We make further functional form assumptions that were left general in the model set up. We assume the utility function is CRRA, $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$, and we assume there are two occupations. We have three parameters that determine the underlying stochastic wage process for each occupation type y: $\mu(y), \sigma(y), \rho(y)$. With two occupations we have 6 distributional parameters. Our empirical results find that the mean and variance of the wages are increasing in y, so we assume that is the case in the simulation.

In table 3 we report the parameter values we use for the simulation. Note that even after conditioning on the persistence, the mean and average are increasing in y.

Table 3: Moments and calibrated parameters

Definition	Parameter	Value	
Risk aversion	γ	2.10	
Borrowing constraint	$\underline{a}$	-80	
Expense shock loss	$\tilde{a}$	2	
Probability expense shock	ϕ	0.07	
Fixed cost of default	κ	0.15	
Disutility of default	ν	2	
		y	
		1	2
AR(1) Wage process	$\mu(y)$	1.06	1.15
	$\sigma(y)$	0.05	0.05
	$\rho(y)$	0.80	0.86

Moments. For the calibration we will estimate the endogenous model parameters by indirect inference using simulated method of moments to match the targeted moments. To inform the estimation of borrowing constraint $\underline{a}$ we target the share of consumers with negative liquid wealth in the US documented from PSID which is 12%. The likelihood ϕ and the magnitude of wealth shock $\tilde{a}$ affect all moments but they are mainly identified by targeting the position of the bottom of kink in the mean and standard deviation of wages.

To discipline the nine parameters from the stochastic wage process, $\mu(y), \sigma(y), \rho(y)$, $y \in \{1, 2, 3\}$, we will match the mean and standard deviation of wages at the following 5% interval percentiles p1, p5, p10,..., p60 of the wealth distribution, of which there are 13.

Results. We now analyze the key features of the steady-state equilibrium. The households in the model face a risk-reward trade-off. When they are hit by the occupation shock, they choose from occupations with different distributions of wages, where the mean and standard deviation of wages across occupations are positively correlated. Workers with higher wealth are better able to self-insure and therefore have a greater aptitude for risk. They choose occupations that offer higher wages, even though the risk of a low draw is higher. In the case of a bad draw, they rely on their wealth to smooth consumption until they get a better draw or until they are presented with a new occupational shock opportunity. When there is occupational choice in addition, households engage in precautionary occupational sorting by choosing lower ranked occupation in order to have less risk exposure.

However, this monotonic relationship between wealth and occupational choice breaks down when workers' wealth becomes sufficiently negative and there is default. Default effectively truncates the

wage distribution below, which implies a higher mean and a lower variance of this truncated distribution. As a result, households that can barely smooth consumption in the presence of risk, gamble for resurrection.

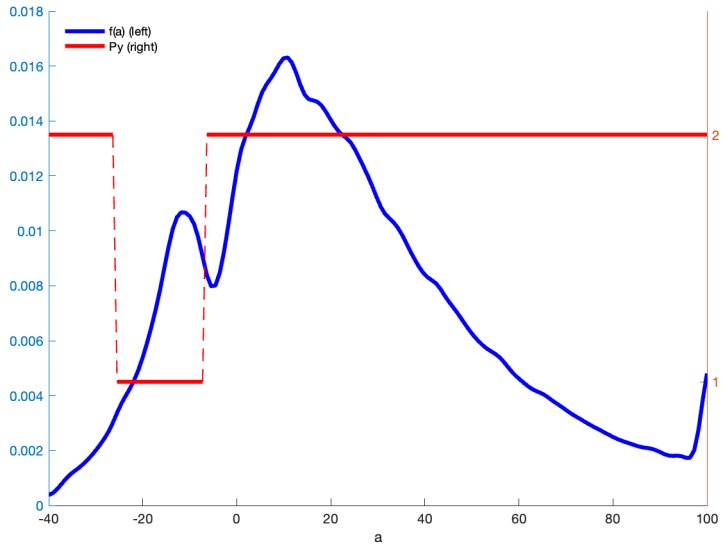


Figure 9: Simulated Economy: Occupational choice and stationary wealth distribution

Figure 9 shows the occupational decisions of workers across different levels of wealth (right axis). Households with positive wealth do not default and their occupational decision is primarily driven by the risk of wage draws. However, for asset levels right below zero, there is no default yet, but to smooth consumption, households choose occupation 1 for minimal exposure to risky wages. When wealth is sufficiently negative, default kicks in and households effectively face the wage distribution that is truncated from below, with higher mean and lower variance.

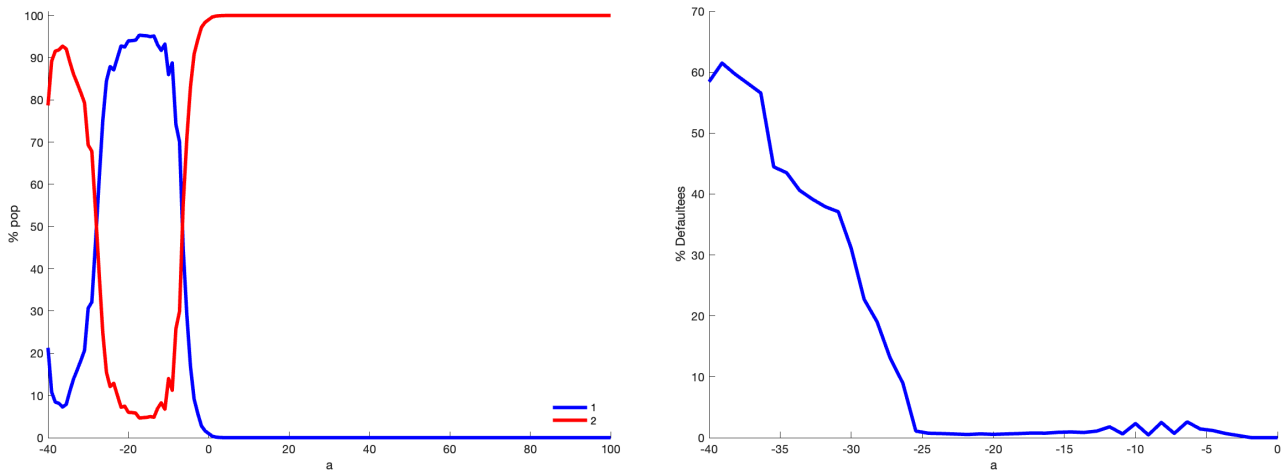
Figure 9 also plots the stationary asset distribution in this simulation.

Figure 10a plots the distribution of occupational choice. Even if the policy decision for occupational choice is deterministic in wealth, the distribution of choices is not because not all households are able to optimally choose a new occupation each period. This happens only with probability λ.

Figure 10b shows that default is highest for those with most debt. The probability of default for those at the bottom of the distribution is 60%, dropping off to zero for those with $250,000 in debt.

Next, we investigate the properties of the wealth distribution in the model in comparison with the data. Table 4 compares three moments of the wealth distribution in the simulated economy with the data. It shows that the simulated economy has wealth levels that are too low compared to the empirical distribution, yet a fraction of households in debt that is too high.

The fact that wages in the simulated economy are below those in the data is also evident when we



(a) Probabilities of Occupational choice

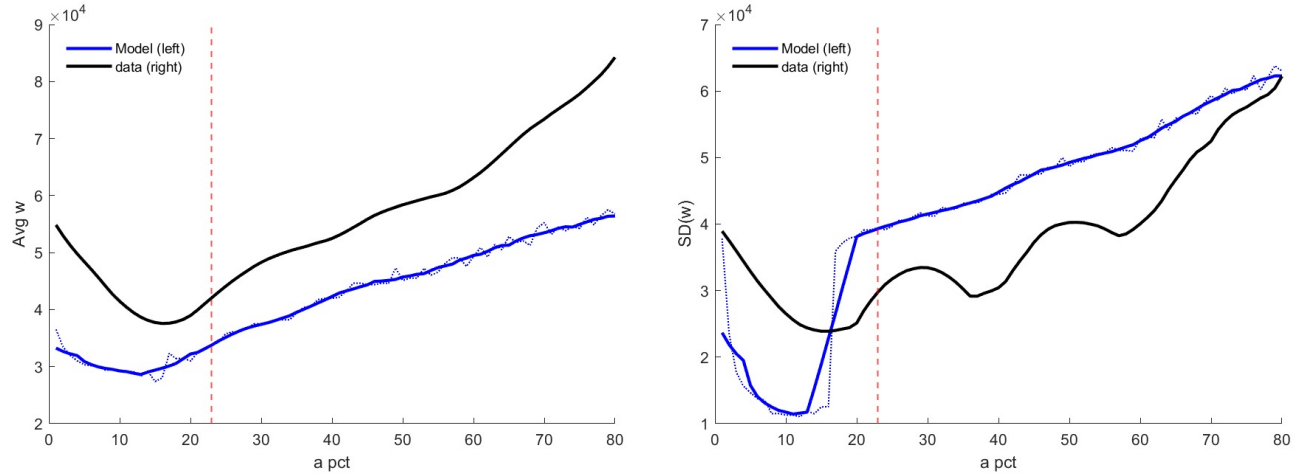
(b) Probability of default on interest payments

Figure 10: Simulated Economy: Probabilities of Occupational Choice and Default

	Data	**Model**
Wage (Median)	\$53k	\$45k
Wealth (Mean)	\$234k	\$222k
Workers with negative wealth (%)	14	23

Table 4: The properties of the wage and wealth distribution in the simulated economy compared to the data (PSID)

plot wages against wealth (see Figure 11a. The model does however capture the V-shaped pattern of wages, with the bottom of the kink around the wealth percentile where we observe it in the data.



(a) Average Wage

(b) Standard deviation of wages

Figure 11: Simulated Economy and Data: Mean and Standard Deviation of Wages

Likewise, the level of the standard deviation in the simulated economy is below that in the data, but there is a V-shaped pattern also in the simulated economy, positioned around where that same pattern is in the data.

Appendix

A Description of Non-US Data

A.1 UK Household Longitudinal Study

Our data are drawn from the UK Household Longitudinal Study (UKHLS), an annual panel survey dataset with information on wages, wealth and demographic characteristics for a representative sample of individuals in the UK. In our main analysis we focus on information from waves 4 and 8 of the survey, collected throughout 2012/13 and 2016/17 respectively, as these waves contain a set of additional detailed questions on wealth and debts.

Wages. Our primary measure of pay is gross labour income (or "wages") in the month before the individual's interview was conducted. Along with usual pay in their main job, this includes additional pay (such as bonuses or overtime), pay from any second jobs, and profits (or losses) from self-employment.[6] Wages are censored at £15,000 a month.[7]

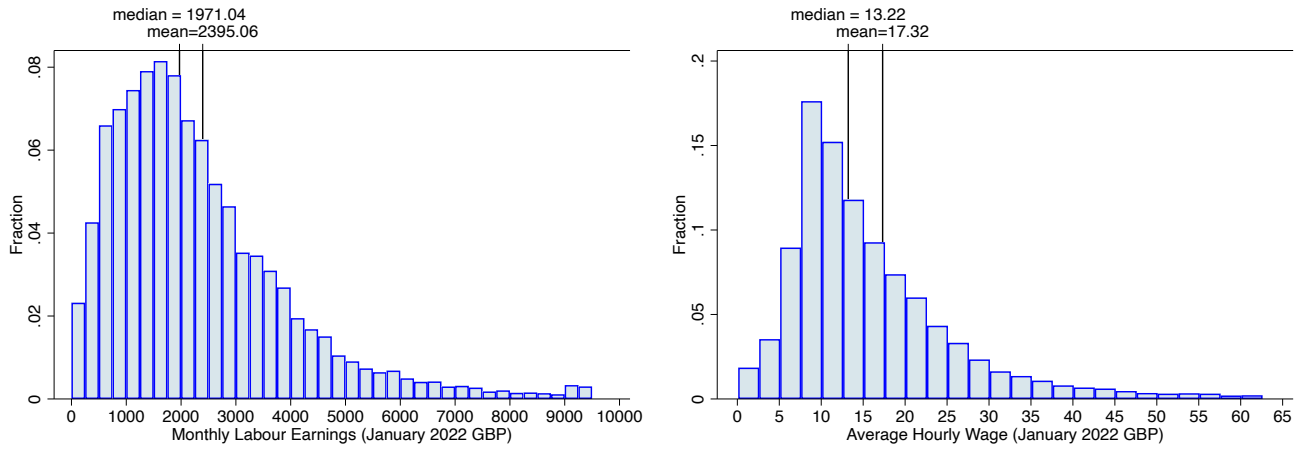
Figure 12 shows the distribution of wages and wages for individuals who were aged 16-65 and in work at the time of their interview, inflated to January 2022 GBP using the Consumer Price Index (CPI). We have dropped any individuals with negative or zero labour income due to self-employment losses or voluntary work. For legibility, we have omitted the top one percent of earners from the histograms – although we have included these high earners in the reported statistics, and keep them for our analysis. Median monthly labour wages is £1971 and the mean is £2395; median hourly wages are £13.22 and the mean is £17.32.

Wealth. We construct two measures of wealth. The first is a measure of net liquid wealth (NLA), defined as the sum of amounts held on deposit or in investment accounts net of any credit card or other non-housing debt (such as personal loans or hire purchase agreements). This is intended to capture wealth the individual can access, or debts they may be liable to repay, at relatively short notice. Our second measure adds to NLA the value of residential housing net of outstanding debt, and so captures individuals' net total wealth (NTA).

We use the following data collected in waves 4 and 8 of UKHLS to construct our measure of net liquid wealth:

[6] As a result of self-employment losses, 0.2% of the sample have negative monthly wages. In practice, we use the logarithm of wages as the dependent variable in our analysis and so drop these individuals from our sample.

[7] Less than 0.25% of our final sample has censored wages.



(a) Monthly Wages

(b) Hourly Wages

Figure 12: Wage distributions, January 2022 GBP

- **Liquid wealth** L**.** Individuals were asked whether they held savings or investments, either in their sole name or jointly with others, in any of (1) a savings or deposit account, (2) national savings account, (3) ISA (cash only) account, (4) ISA (investment) account, (5) premium bonds or (6) other type of account. For each of these six types of account an individual reported holding, they were asked how much they held in total across all account of that type. We have taken the sum of savings held across all six account types as a measure of liquid wealth.

- **Credit card debt** C_d**.** Individuals were asked whether they held any store-cards or credit cards, either in their sole name or jointly with others. If so, they were asked to report their credit limit and outstanding debts. We have calculated the individual's credit card debt C_d as the sum of outstanding debts across all reported cards.

- **Other debt** D**.** Individuals were asked whether they held any other debts, excluding mortgages, either in their sole name or jointly with others. The categories of debt included in this measure are: hire purchase agreements, personal loans (e.g. from a bank), mail-order purchase agreements, DWP social loan funds, any loans from a private individual, overdrafts, student loans, or "other". We have calculated D as the sum of debts across all possible categories.

Our definition of net liquid wealth is

$$NLA = L - C_d - D \tag{9}$$

We note two features of the data. First, individuals sometimes do not know how much they held in savings or owed in debt. In these cases, individuals were asked a series of branching yes/no questions

to attempt to bound the value. For example, if an individual did not know how much they held in a savings account, they were asked whether it was more than £1000. If yes, they were then asked whether it was more than £5000; if no, they were asked whether it was more than £500 (and so on). The actual values used in these questions were different for each type of savings account, credit card or other debt.

The results of these branching questions allow us to place the value of wealth or debts into one of six intervals when the individual does not know the precise value. We have recorded the wealth of an individual in interval $[a, b)$ as $\frac{a+b}{2}$. Note that, if an individual reports that their wealth are higher than V_5, we can only place their asset holdings in interval $[V_5, \infty)$. For these individuals, we have assumed that their wealth fall in interval $[V_5, 2V_5)$.

Second, some wealth or debts are reported as being held jointly with others. If an asset is entirely held with others, we have divided the value by the total number of joint holders. For example, a savings account containing £1000 held jointly with a spouse is recorded as £500. If only part of an asset or liability is held jointly, with the rest held entirely in the individual's sole name, the individual is asked to estimate the value of the asset they hold in their sole name. For an asset of value A, of which $S \leq A$ is held solely and the remainder is held jointly with $N - 1$ other people, we have recorded the value of the asset for the individual as $S + \frac{A-S}{N}$. Some individuals who report that an asset is in part held jointly are unable to estimate the amount they hold individually S. In these cases, we have assumed that the asset is entirely held by the individual (i.e. $S = A$).

We have also used household-level data on net housing wealth for our total wealth measure. This information is, in principle, collected in every wave of UKHLS — although the questions are only asked to households (rather than individuals), and then only when the household enters the sample or has moved since the last interview. We therefore collect information on housing wealth from the full UKHLS (and BHPS, the predecessor survey) panel.

- **Housing wealth H_a.** Households which report owning their home outright, owning their home with a mortgage, or being part of a shared ownership scheme (part-ownership, part-rent) are two sets of questions we could use to estimate the value of their current housing wealth:

  1. Households are asked to estimate the current market value of their home in each wave.
  2. They are also asked when they bought their home, and how much they paid for it, when they join the sample or when they move house.

  We currently have used the household's own estimate of the current market value of their property, since this is likely a better reflection of current wealth for individuals who have lived in their house for a number of years.

- **Housing debt H_d.** We have combined two measures of housing debt to maximise the available sample.

    1. Individuals who do not own their home outright are asked (when they join the sample or when they move house) how much they borrowed to buy their house (M), when they borrowed it (T_{buy}), and the number of years remaining on their mortgage (T_{left}). We have assumed that households will run down their housing debt linearly, and so estimated the remaining housing debt for a household interviewed in year T as:

    $$H_d = M \times \left(\frac{T_{left}}{T + T_{left} - T_{buy}} \right)$$

    For individuals who own their home outright, we have set $H_d = 0$.

    2. In wave 1 of UKHLS, individuals were asked to provide an estimate of the total debt outstanding on their homes. While the timing of this estimate does not coincide exactly with the timing of other questions (as it is from wave 1 and not waves 4 or 8) it provides the most recently available estimate of outstanding housing debt for all individuals who have not moved since wave 1 of UKHLS.

    We have combined the two measures to create a measure of housing debt which is non-missing for as many observations as possible. We prioritised information from our second measure of debt and, for individuals whose value of debt from this measure is missing, used the first measure of estimated outstanding debt where possible. Of the 18,349 individuals in wave 4 of UKHLS who own their home with a mortgage:

    – 12,830 (70%) individuals have missing values for the first measure of housing debt;
    – 9,199 (50%) individuals have missing values for the second measure of housing debt;
    – 4,409 (24%) individuals have missing values for our combined measure of housing debt.

Our measure of net total wealth is

$$NTA = L - C_d - D + H_a - H_d \tag{10}$$

Figure **??** shows the distribution of each of these measures, again inflated to January 2022 GBP using the CPI. As with the pay distribution in Figure 12, we have omitted the tails of each distribution from the histograms for readability – in this case, we have omitted the top and bottom five percent – and the mass at precisely zero for each measure. Both distributions have a mass point at zero: 8.4 percent

Table 5: Sample Composition

	Wave 4	Wave 8	Pooled
Whole sample	47071	39288	86359
of which			
(i) working or self-employed	25966	22053	48019
(ii) age 16-65	25339	21442	46781
(iii) positive earnings	24829	20781	45610
(iv) non-missing total assets	18446	15527	33973

of the sample report exactly zero net total wealth, and 18.4 percent report zero net liquid wealth. Both distributions are positively skewed; to reduce the influence of individuals with extreme asset values, we express wealth as percentiles in our analysis.[8]

Sample. Our main sample contains 33,973 individuals. Of the 86,359 individuals interviewed in UKHLS waves 4 or 8, we kept only the 48,019 who were working or in self-employment. Of these, we kept only the 46,781 aged between 16 and 65; then the 45,610 of these with positive wages, and finally the 33,973 of these for which all variables required to calculate net total wealth are non-missing. Table 5 shows how these sample selection criteria affect individuals from each wave separately, together with the pooled sample which forms the basis of our analysis.

B Robustness of Empirical Results

In this Appendix, we collect together analyses to demonstrate the robustness of our main empirical facts. In particular, we show that the qualitative features of our empirical results are not sensitive to: (i) a specific definition of wealth; (ii) specific measures of wages; or (iii) the particular restrictions we place on our sample, including the period we study or the age or composition of households in our sample, or features of the earnings of households with precisely zero wealth.

B.1 Alternative definitions of wealth

The measure of wealth we use throughout our analysis is *total* wealth—that is, the sum of cash, bonds, stocks, the value of any business, pension funds, housing, and other real estate, and vehicles — net of any mortgage or other debt. In this section, we consider a more restrictive measure of wealth—*liquid* wealth—which is intended to capture resources that are easier for households to access or borrow

[8]This means, for example, that an individual is assigned 0.5 if their asset holding is equal to the median.

against on short notice. Liquid wealth excludes the value of home equity (i.e., the net value of a primary residence after subtracting any mortgage), the value of vehicles, and the value of pension funds.

Figure 13 shows estimates of the non-parametric relationship from equation (1) between the percentile rank of liquid wealth and wages. There is also a clear and precisely estimated non-monotonicity when using this alternative definition of wealth. We note that the inflection point in the relationship is further to the right in the distribution of liquid wealth than is the case for total wealth, but still corresponds to the point at which liquid wealth is zero.

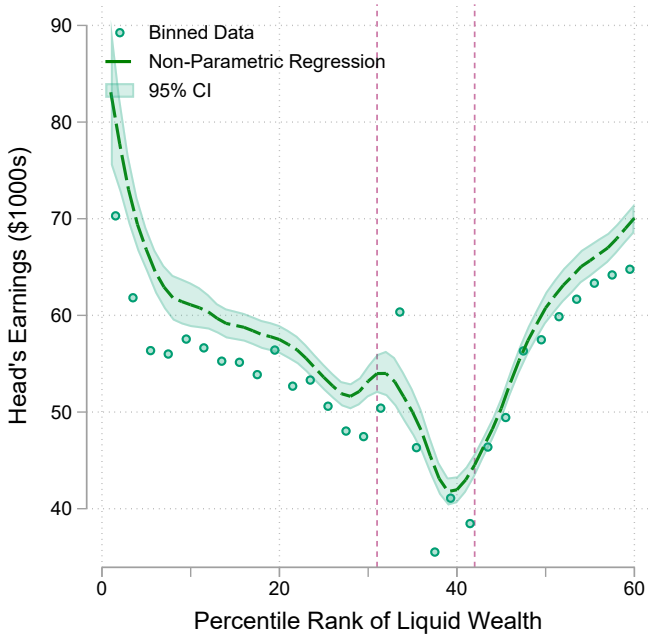


Figure 13: Non-monotone relationship between *liquid* wealth and annual wages

Notes. This is the estimated kernel regression in equation (1), along with 95 percent confidence intervals and a binned scatter plot of the raw data where each point represents average wages in a two-percentile interval on the wealth distribution. The measure of wealth is *liquid* wealth, i.e. the sum of cash, bonds, stocks, the value of any business, net of any non-housing, non-vehicle debt; or, to put it differently, total wealth as defined in Section 2.1, minus the value of pension funds, home equity and vehicles.

B.2 Alternative definitions of wages

The measure of wage in our main analysis is the household head's annual labor earnings, expressed in dollars per year. Figure 14 shows that the main result is robust to alternative definitions of the wage. In Panel (a), we use the percentile rank of the household head's annual labor earnings, which reduces the influence of individuals with extremely high or low wages; in Panel (b), we construct a measure of permanent income from the residual of a Mincer equation $\ln(y_{i,t}) = \alpha + \sum_{k=1}^{3} \beta_k age_{i,t}^k + \phi_t + \epsilon_{i,t}$ to

Table 6: Robustness to Sample Composition

	(1) Baseline	(2) Include 0s	(3) Include Single Men	(4) Include Single Women	(5) All Ages
Wealth $(a_{i,t} \geq 0)$	1281.147***	1269.019***	1289.909***	1185.963***	1158.004***
	(121.330)	(117.142)	(97.061)	(78.876)	(75.611)
Wealth $(a_{i,t} < 0)$	-981.798***	-947.489***	-864.5348***	-845.5725***	-888.1525 ***
	(333.278)	(331.541)	(218.616)	(143.390)	(156.922)
Controls					
Age	Yes	Yes	Yes	Yes	Yes
Race	Yes	Yes	Yes	Yes	Yes
Education	Yes	Yes	Yes	Yes	Yes
N	26,317	27,699	35,544	48,411	64,313

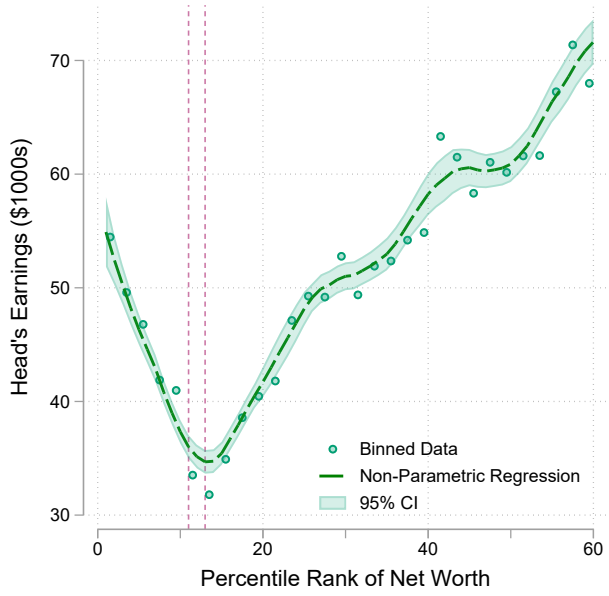
Notes. In all regressions we additionally control for state and year fixed effects. Standard errors clustered at the household level. * $p < 0.10$, ** $p < 0.5$, *** $p < 0.01$

control for the age-profile of earnings, and use the percentile rank of this permanent income as our measure of wage; in Panel (c) we use the percentile rank of income net of taxes, and including transfers from major welfare programmes in the US (i.e. the EITC and food stamps); and in Panel (c) we use the percentile rank of the hourly (rather than annual) wage. The non-monotonic relationship we estimate is qualitatively—and, indeed, quantitatively—robust to these alternative measures of the wage.

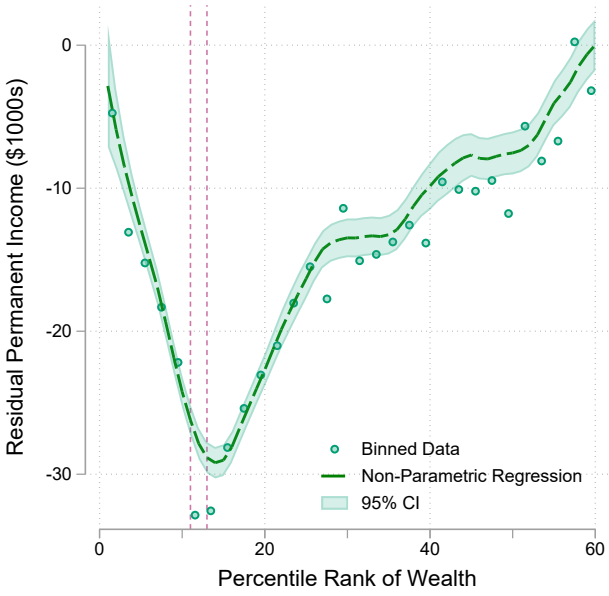
B.3 Alternative sample restrictions

In this section, we demonstrate the robustness of our first empirical fact to different sample selection choices. Table 6 contains estimates of equation (2) separately for different samples. The "Baseline" estimates in column (1) are therefore identical to column (3) of Table 1 in the main text. We show that this main result is qualitatively and quantitatively similar if we: include those who do not work (and so have a wage of zero) in column (2); if we include single men along with married household heads in our sample in column (3); if we include single women along with married household heads in our sample in column (4); and if we include household heads younger than 25 or older than 55 in column (5).

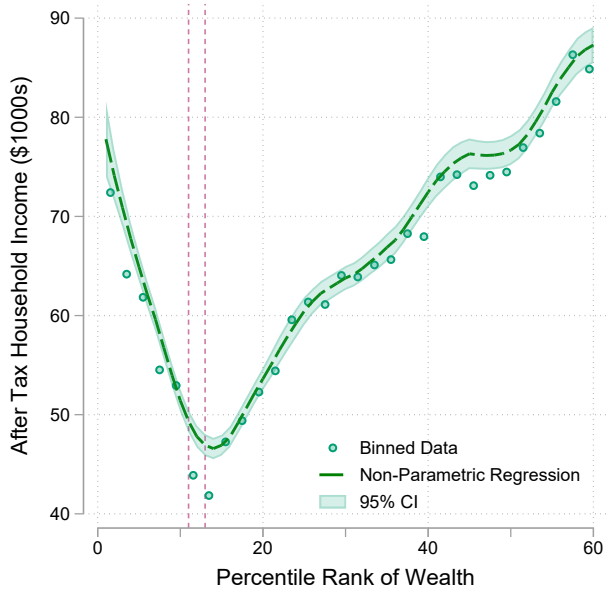
A separate concern is that the non-monotonicity we document may be driven by some particular feature of the labour market behavior of those with precisely zero wealth, rather than a more general finding. We demonstrate that this is not the case in Table 7, which contains estimates of equation (2) excluding individuals located in increasingly large regions around zero wealth. The main feature is



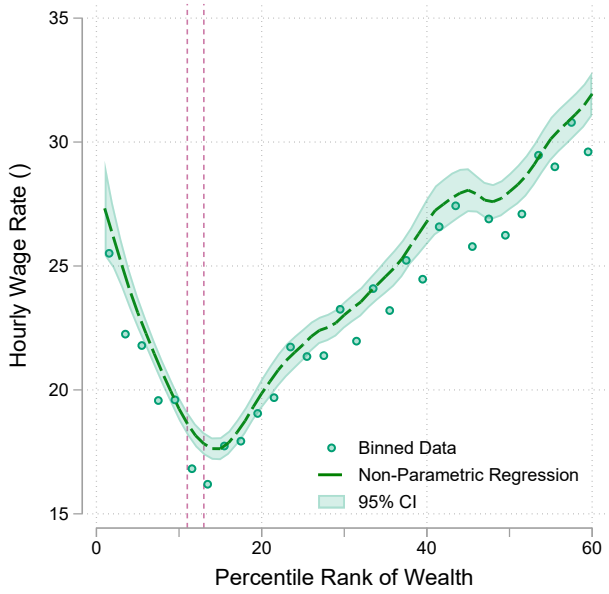
(a) Annual Wage



(b) Residual Permanent Income



(c) Annual Post-Tax Household Wage



(d) ourly Wage

Figure 14: Non-monotone relationship between wealth and various definitions of annual wages

Notes. Estimated kernel regressions from equation (1), along with 95 percent confidence intervals and a binned scatter plot of the raw data where each point represents average wages in a two-percentile interval on the wealth distribution, for various definitions of the wage. Panel (a) uses the percentile rank of the annual wage (rather than its level); Panel (b) uses the percentile rank of permanent income, defined as the residual from Mincer regression $w_{i,t} = \alpha + \sum_{k=1}^{3} \beta_k age_{i,t}^k + \phi_t + \epsilon_{i,t}$; Panel (c) uses the percentile rank of post-tax (rather than gross) annual wages; and Panel (d) uses the hourly (rather than annual) wage.

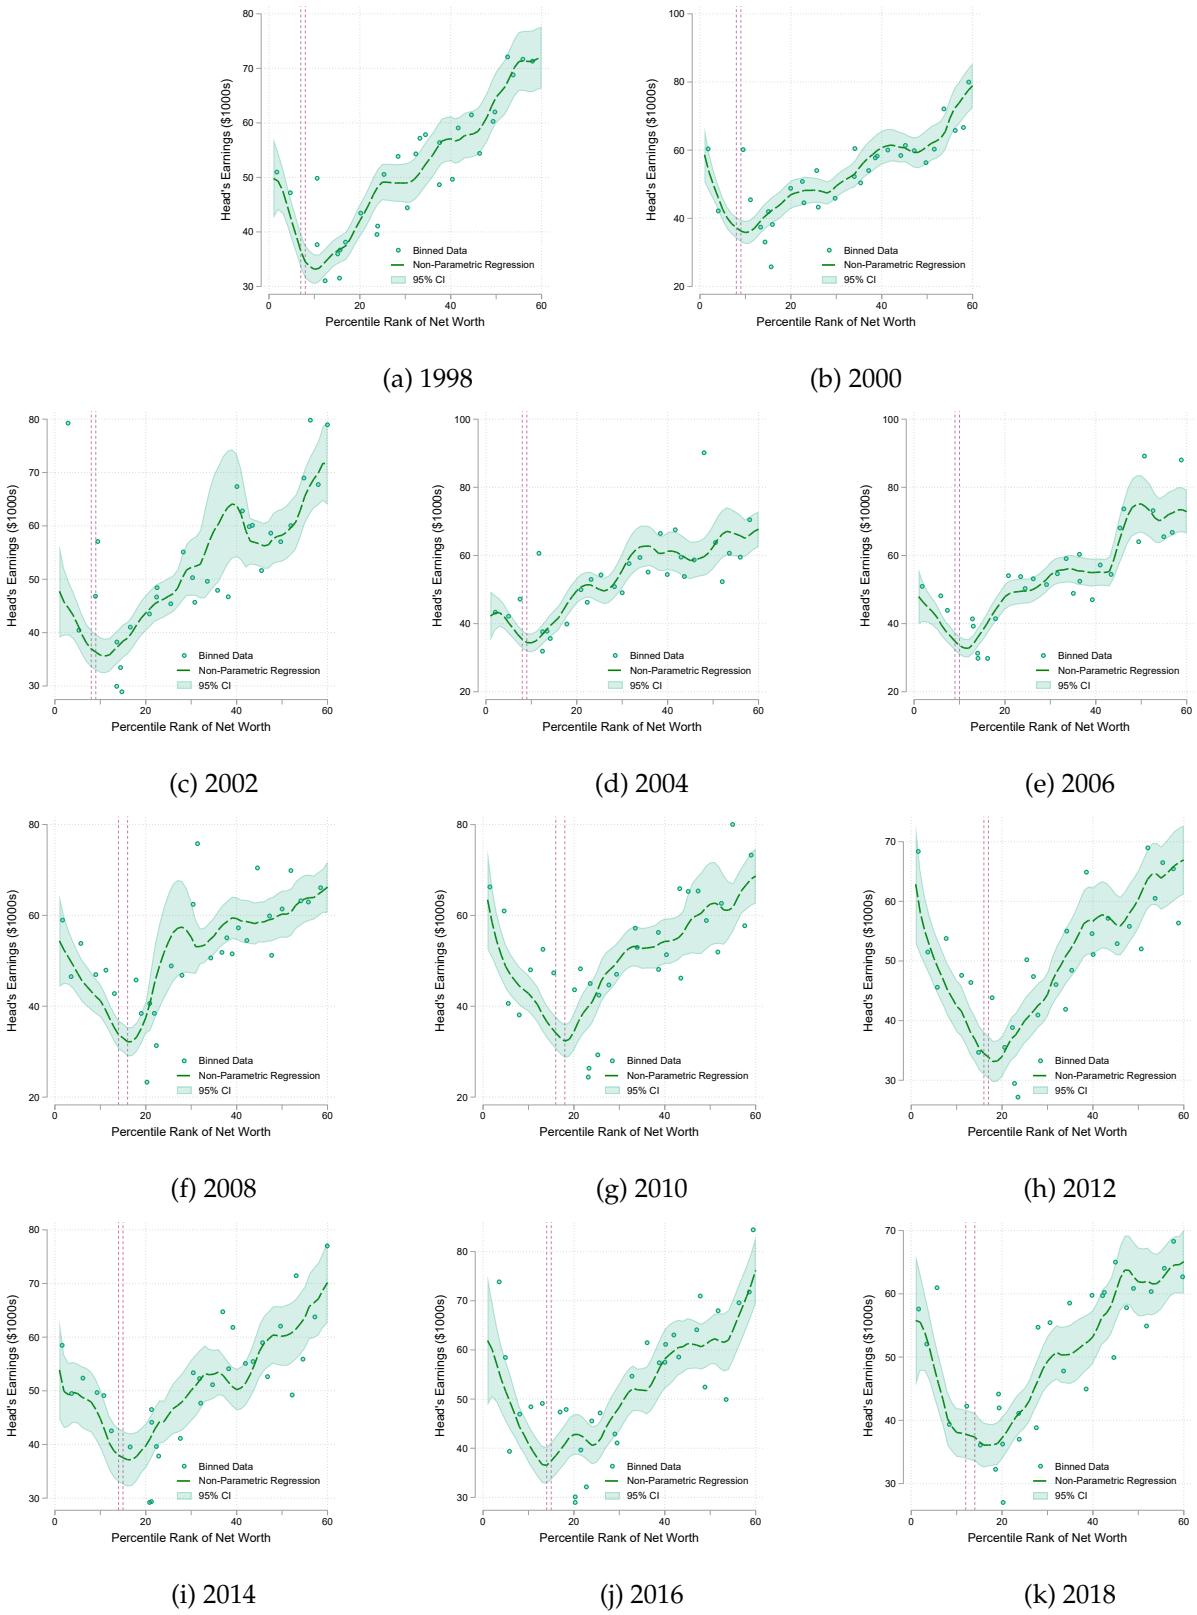
Table 7: Parametric tests for non-monotonicity, excluding regions around zero wealth

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Wealth ($a_{i,t} \geq 0$)	1281.147***	1299.718***	1298.916***	1301.256***	1306.777***	1328.360***	1362.924***	1417.841***
	(121.330)	(124.067)	(124.042)	(124.595)	(125.574)	(129.452)	(135.383)	(145.496)
Wealth ($a_{i,t} < 0$)	-981.7983***	-959.0422 ***	-945.4169***	-886.7014***	-913.653***	-877.0079**	-1037.308**	-903.009***
	(333.278)	(335.848)	(337.412)	(318.423)	(328.156)	(360.219)	(411.704)	(515.039)
Excluded Region	n/a	[-100,100]	[-250,250]	[-500,500]	[-1k,1k]	[-2.5k,2.5k]	[-5k,5k]	[-10k,10k]
Controls								
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Race	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Education	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	26,317	25,816	25,772	25,688	25,502	24,930	24,068	22,692

Notes. In all regressions we additionally control for state and year fixed effects. Standard errors clustered at the household level. * $p < 0.10$, ** $p < 0.5$, *** $p < 0.01$

robust across all specifications, which exclude individuals from a region of up to $10,000 either side of zero wealth.

Finally, in Figure 15, we show our main non-monotonicity separately by PSID wave. While the estimates are less precise due to considerably smaller sample sizes in each wave, the non-monotonicity is present in all years.



(a) 1998 (b) 2000

(c) 2002 (d) 2004 (e) 2006

(f) 2008 (g) 2010 (h) 2012

(i) 2014 (j) 2016 (k) 2018

Figure 15: Non-monotonicity estimated separately by year

References

AIYAGARI, S. R. (1994): "Uninsured Idiosyncratic Risk and Aggregate Saving," *The Quarterly Journal of Economics*, Vol. 109, No. 3., 659–684.

BALEY, I., A. FIGUEIREDO, C. MANTOVANI, AND A. SEPAHSALARI (2022): "Self-Insurance and Welfare in Turbulent Labor Markets," Tech. rep., Working Paper.

BEWLEY, T. F. (1980): "Interest Bearing Money and the Equilibrium Stock of Capital," Mimeo.

BLUNDELL, R., L. PISTAFERRI, AND I. SAPORTA-EKSTEN (2016): "Consumption Inequality and Family Labor Supply," *The American Economic Review*, 106, 387–435.

CHATTERJEE, S., D. CORBAE, M. NAKAJIMA, AND J.-V. RÍOS-RULL (2007): "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default," *Econometrica*, 75.

DAMINATO, C. AND L. PISTAFERRI (2023): "Heterogeneity in returns to wealth and consumption inequality," .

EECKHOUT, J. AND I. LINDENLAUB (2019): "Unemployment Cycles," *American Economic Journal: Macroeconomics*.

EECKHOUT, J. AND A. SEPAHSALARI (2024): "The effect of wealth on worker productivity," *Review of Economic Studies*, 91, 1584–1633.

FAN, J. AND Q. YAO (1998): "Efficient estimation of conditional variance functions in stochastic regression," *Biometrika*, 85, 645–660.

GOLLIER, C., P.-F. KOEHL, AND J.-C. ROCHET (1997): "Risk-taking behavior with limited liability and risk aversion," *Journal of Risk and Insurance*, 347–370.

GOURINCHAS, P.-O. AND J. A. PARKER (2002): "Consumption over the life cycle," *Econometrica*, 70, 47–89.

GROES, F., P. KIRCHER, AND I. MANOVSKII (2015): "The U-shapes of occupational mobility," *The Review of Economic Studies*, 82, 659–692.

GUVENEN, F., F. KARAHAN, S. OZKAN, AND J. SONG (2021): "What do data on millions of US workers reveal about lifecycle earnings dynamics?" *Econometrica*, 89, 2303–2339.

HUANG, J. E. AND X. QIU (2022): "Precautionary mismatch," Tech. rep., Mimeo, University of Pennsylvania.

HUGGETT, M. (1993): "The risk-free rate in heterogeneous-agent incomplete-insurance economies," *Journal of Economic Dynamics and Control*, 17, 953–969.

KOŞAR, G., D. MELCANGI, L. PILOSSOPH, AND D. G. WICZER (2023): "Stimulus through Insurance: The Marginal Propensity to Repay Debt," .

KRUSELL, P., J. LOU, AND V. J. RÍOS-RULL (2019): "Wealth, Wages, and Employment," *Working Paper*.

LEWIS, D. J., D. MELCANGI, AND L. PILOSSOPH (2019): "Latent heterogeneity in the marginal propensity to consume," *FRB of New York Staff Report*.

SINN, H.-W. (1982): "Kinked utility and the demand for human wealth and liability insurance," *European Economic Review*, 17, 149–162.

STIGLITZ, J. E. AND A. WEISS (1981): "Credit rationing in markets with imperfect information," *The American economic review*, 71, 393–410.