Labor Market Polarization with Hand-to-Mouth Households*

Johannes Wacks †

February 11, 2024

Abstract

I argue that borrowing constraints are important for understanding the welfare and output consequences of labor market polarization. I build a general equilibrium incomplete markets model featuring three occupations, a continuum of skill types, occupation-specific human capital and a realistic share of hand-to-mouth households. A fall in the price of capital causes the routine wage to decline relative to wages in the other occupations. The presence of a large share of households close to the borrowing constraint inefficiently protracts the reallocation of labor away from the declining routine occupation. Policies that alleviate the borrowing constraint upon switching the occupation raise social welfare and output, speeding up the reallocation of workers away from the declining routine occupation. While disadvantageous for the high-skilled, the policies benefit medium- and low-skilled workers.

JEL Codes: D31, E21, E24, J24, J62

Keywords: Wealth distribution, Technological change, Labor market reallocation, Borrowing constraints

*I thank Klaus Adam, Pedro Amaral, Mark Bils, Elizabeth Caucutt, Juan Carlos Conesa, Matias Cortes, Andreas Gulyas, Tom Krebs, Alexander Matusche, Matthias Meier, José-Víctor Ríos-Rull, Ayşegül Şahin, Frederick Zadow, as well as participants at numerous seminars and conferences for helpful comments and suggestions. I gratefully acknowledge funding by the Stiftung Geld und Währung and by the Graduate School of Economic and Social Sciences at the University of Mannheim. The views expressed in this paper are my own. They do not necessarily coincide with the views of the Deutsche Bundesbank or the Eurosystem.

†University of Mannheim; Deutsche Bundesbank; Wilhelm-Epstein-Strasse 14, 60431 Frankfurt am Main, Germany; E-mail: johannes.wacks@bundesbank.de

1 Introduction

Technological change has had a major impact on labor markets in recent decades. Many jobs previously performed by humans have become automated and are now completed by machines. This trend has benefited workers in some occupations while harming those in others. Consequently, wages and employment shares in occupations intensive in routine, i.e. easily codifiable tasks (e.g. assembly line workers, bookkeepers) tasks have declined in recent decades compared to those in manual (e.g. taking care of the elderly or children) and abstract occupations (e.g. teachers, managers).[1]

In this paper, I argue that borrowing constraints and the ability to smooth consumption are important to understand the output and welfare consequences of this technological change and the labor market reallocation it has caused. To share in the benefits of technological growth, many routine workers have left their old jobs and switched into occupations in which wages have been rising (Cortes, 2016). These switches, while providing benefits in terms of higher wages in the medium to long term, were often accompanied by initial wage losses. This pattern, which can be rationalized by the presence of occupation-specific human capital, makes occupational choice a dynamic investment decision. The ability to smooth consumption, i.e. the distance from the borrowing constraint, is therefore an important determinant of whether a worker decides to switch or stay in her old job.

The main contribution of this paper is to embed routine-biased technological growth into a general equilibrium model of the US economy featuring incomplete markets and occupation-specific human capital. I use the model to demonstrate that borrowing constraints play an important role for the reallocation of labor away from routine occupations. The aggregate labor market transition is impeded by the presence of "hand-to-mouth" households, i.e. those with few or no liquid assets. The closeness to the borrowing constraint makes them unwilling to invest in future earnings growth at the cost of short-term wage losses. This gives rise to an inefficiency, as the discounted gains from switching can outweigh the initial costs. In short, hand-to-mouth workers accumulate human capital in the declining occupations for too long. Building on this insight, I use the model to show how government policies aimed at alleviating the borrowing constraint of occupational switchers can improve both output and welfare.

To assess whether the friction I propose is empirically relevant, I build on and extend previous empirical findings. Using panel data from the Panel

[1] These empirical patterns have been documented, among others, by Acemoglu and Autor (2011), Autor and Dorn (2013), Autor, Levy, et al. (2003), Cortes, Jaimovich, and Siu (2017), and Goos et al. (2014).

Study of Income Dynamics (PSID), Cortes (2016) shows that, compared to those workers who stayed in routine occupations, workers who switched to abstract or manual occupations saw faster wage growth and thus enjoyed long-term wage gains. Adding to this, I show that a large share of switchers did, however, suffer wage losses in the short run, i.e. in the year immediately following the switch. Next, I use data from the Survey of Consumer Finances (SCF) to demonstrate that a third of routine workers fall under common definitions of being hand-to-mouth, i.e. hold very few liquid assets. This shows that potential short-term wage losses might have posed an important obstacle to an occupational switch for a large fraction of routine workers. Lastly, I bring these two pieces of information together, turning to recent waves of the PSID, which provide information on asset holdings. I ask whether being hand-to-mouth is predictive of leaving the routine occupations for a new job, and find that, in line with the mechanism described above, hand-to-mouth households are less likely to make such a switch.

I then build a heterogeneous agent model with uninsurable idiosyncratic labor income risk, that takes account of these empirical findings. On the supply side, a representative firm uses two types of capital (Information and Communication Technology (ICT) and non-ICT) and three types of labor (manual, routine and abstract labor) as inputs. The driving force of technological change is an exogenous fall in the price of ICT capital. On the demand side, the model features households that die stochastically and are then replaced by newborns (perpetual youth). Households are heterogeneous in several ways. They have a fixed skill type, which influences their optimal occupational choice. Households accumulate occupation-specific human capital and they are exposed to uninsurable idiosyncratic labor income risk, causing them to engage in precautionary saving. Households save in two assets, a liquid and an illiquid one. This generates a share of hand-to-mouth agents that matches the data, while at the same time allowing me to model a production economy.

I calibrate the model to the US economy, targeting the employment shares in each occupation and the share of income accruing to labor in 1980 and 2020, the share of hand-to-mouth households, as well as the average wage changes of routine workers who switch to either the abstract or to the manual occupation. I solve for the steady state when the price of ICT capital is relatively high (representing 1980) and relatively low (representing today), and compute the perfect foresight transition path between the two. The calibrated production function implies that ICT capital is substitutable with routine work and complementary to manual and especially abstract work. This leads to a lower routine employment share in the new steady state.

Consistent with the empirical evidence and the proposed mechanism the model predicts an important role of liquid assets and hence the ability to smooth consumption for the switching behavior of routine workers. I demonstrate this, first, by zooming in on individual policy functions, showing that routine workers leave the occupation earlier if they possess liquid asset savings than if they are hand-to-mouth. Second, in counterfactual simulations of the economy in which all households choose their occupation like the ones who are rich in liquid assets, the employment share in the abstract occupation would have been two to three percentage points higher along the transition than in the baseline. In the appendix, I further conduct the same regressions of hand-to-mouth status on future switching decisions as before in the PSID, now in a synthetic panel of households simulated from the model. I find the same negative association as in the data.

Next I ask how the policymaker can address the inefficiency introduced by the borrowing constraint by studying policies that target workers who leave the routine occupation. First, the government pays a recurrent transfer to the switchers, partly covering their temporary wage loss. Second, the government hands out a loan to switchers, providing them with additional liquidity upon leaving the routine occupation. Both policies are financed via distortionary labor income taxation.

Focusing on policies targeting the routine to abstract switchers, I find that both programs increase aggregate welfare vis-à-vis the baseline transition without policy. This is because eligible households directly benefit from the programs, and because of general equilibrium effects. The constrained optimal policies lead to a more efficient allocation of the labor force across occupations, and hence output and average labor productivity increase. The majority of the population enjoys wage gains: as more households switch to the abstract occupation with the policies in place, wages of abstract workers decline while wages of manual and routine workers rise. The clear winners of the policies are medium-skilled workers, who are predominantly employed in the routine occupation. Their expected lifetime consumption increases by 1.0% (0.3%) under the optimal transfer (loan) program. Importantly, the transfer program leads to larger efficiency and welfare gains than the loan program because it conditions payments on workers staying in the new occupation. This leads to a more efficient selection of workers into the program.

I validate the robustness of the policy results in a number of ways. Among others, I show that allowing for heterogeneity in preferences does not alter the conclusions and that departing from the assumption of perfect foresight regarding technological change even strengthens the quantitative results.

I discuss next the related literature. In Section 2 I use a small, highly stylized model of dynamic occupational choice to develop the intuition for the key mechanism. In Section 3 I present empirical evidence that motivates and supports the quantitative model. The full quantitative model is presented in Section 4. In Section 5 I study the transition between two steady states, while in Section 6 I conduct the policy experiments. Section 7 concludes.

Related Literature This study relates to three strands of the literature. The most closely related literature embeds labor market polarization into quantitative macroeconomic models. vom Lehn (2020) builds a model in which households are heterogeneous in skill type and endogenously sort into the three broad occupational groups. While in vom Lehn (2020) the household sector can be summarized by a representative agent, in my model workers face uninsurable idiosyncratic income risk (Aiyagari, 1994; Bewley, 1983; Huggett, 1993).[2] Importantly, I also add occupation-specific human capital, which makes occupational choice an investment decision.

Moll et al. (2022) add household heterogeneity in skills and dissipation shocks to wealth accumulation to the task-based framework of Acemoglu and Restrepo (2018). They simulate a trend in automation and find that, by driving up the interest rate on capital, it leads to higher wealth inequality. I differ in modelling uninsurable idiosyncratic income risk and occupation-specific human capital, and by focusing on the normative aspects of labor market polarization. The interest in policy analysis is shared by Jaimovich, Saporta-Eksten, et al. (2021). However, while they put the focus on labor market frictions, I zoom in on the importance of borrowing constraints.

Beraja and Zorzi (2022) (BZ) study a model of endogenous automation. When making their automation decision, firms impose a negative externality on workers as they do not take into account the workers' borrowing constraints and costly reallocation process. Hence, automation can be inefficiently high. My model differs in several ways from BZ, offering at least three important additional insights. First, following a large literature, I model labor market polarization as exogenous, routine-biased technological change that hereteterogeneously affects different worker skill types. I can thus speak to the heterogeneous welfare consequences of labor market policies along the skill distribution. Second, in contrast to BZ, I explicitly

[2] Kikuchi and Kitao (2020) study labor market polarization in an incomplete markets model, but in a partial equilibrium setting. General equilibrium effects are important when I study labor market policies.

model occupation-specific human capital. This makes occupational choice an investment decision and on aggregate the reallocation of workers across occupations becomes inefficiently slow. Hence, while BZ emphasize the role of firms and their automation decision in determining the pace of aggregate reallocation, I emphasize the role of households and their dynamic occupational choice. This also leads to a slightly different view on (constrained) optimal policy: while BZ's framework gives rise to the notion that the government should slow down automation, e.g. by raising an automation tax, in my framework with occupation-specific human capital the government should speed up reallocation by alleviating borrowing constraints, thereby raising productive efficiency. Third, my model features a much richer household side, as I rely on the two-asset specification of Kaplan, Moll, et al. (2018). This enables me to show, for instance, that the highlighted friction is not only highly relevant for poor but also for wealthy hand-to-mouth households, who are far greater in number. Moreover, I supplement my model with novel empirical evidence on the interplay between liquid assets and occupational choice.

The second related strand of the literature models occupation-specific human capital and wealth accumulation simultaneously. That human capital is at least in part tied to a worker's specific occupation or task has been documented in Cortes and Gallipoli (2018), Gathmann and Schönberg (2010), Kambourov and Manovskii (2009b), Sullivan (2010), and Traiberman (2019). Efforts to study occupational choices and wealth accumulation jointly go back at least to Bernhardt and Backus (1990), Evans and Jovanovic (1989), and Galor and Zeira (1993). A key takeaway from this literature is that when returns to tenure differ by occupation, borrowing constraints can prevent young workers from choosing an occupation that features a steep earnings profile and would maximize their lifetime income. While I also allow for occupation-specific returns to tenure, I show that the differential wage growth across occupations offers a rationale for labor market policies supporting occupational switches even of experienced workers.

The third related strand is the empirical literature exploring labor market polarization (Acemoglu and Autor, 2011; Autor and Dorn, 2013; Cortes, 2016; Cortes, Jaimovich, Nekarda, et al., 2020). I add novel empirical evidence to this in Section 3.

2 A Tractable Model of Dynamic Occupational Choice with Borrowing Constraints

Here I illustrate the core mechanism at work in the full quantitative model of Section 4 in an analytically tractable environment. When relative wages between occupations change over time, a worker potentially makes differing occupational choices depending on whether she is at the borrowing constraint or not. A necessary model ingredient is that human capital is occupation-specific, as in this case occupational choice becomes an investment decision. A trade-off arises between smoothing consumption and building up human capital in the occupation whose wages are growing.

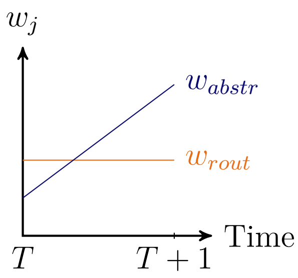

Figure 1: Time path of wages per efficiency unit of labor in routine and abstract occupation

The model has two time periods, T and $T+1$. There are two occupations, routine and abstract. The wage per efficiency unit of labor in the routine occupation is one in both periods, i.e. $w_{r,T} = w_{r,T+1} = 1$. The abstract wages are $w_{a,T} = \frac{1}{\omega}$ and $w_{a,T+1} = \omega$, with $\omega > 1$ (Figure 1).[3]

Consider a household who makes an occupational choice at the beginning of period T and again at the beginning of period $T+1$. Her income in every period is $y_t = w_{j,t} \cdot h_t$, where $h_t \in \{\underline{h}, \bar{h}\}$ is occupation-specific human capital, with $0 < \underline{h} < \bar{h}$. If the household works in the same occupation j in T and $T+1$, she is an experienced worker in occupation j in $T+1$ with certainty ($h_{T+1} = \bar{h}$). An occupational switch at the beginning of period t leads to full depreciation of human capital ($h_t = \underline{h}$). The household values consumption, with $u(c) = \ln(c)$. I assume that the household arrives with zero assets in T and that she discounts future utility by a factor $\beta \in (0, 1)$.

Suppose that at the beginning of period T the household has not yet

[3]Empirically, of course, abstract workers earn higher wages on average than routine workers, and have done so not only recently but also in the 1980s. However, for certain middle-income workers the potential wages in the two occupations have plausibly followed the pattern shown here. For instance, for some workers it paid more to have a job at an assembly line or to be a bookkeeper than to be a manager or a teacher in the 1980s, while in the 1990s or 2000s this ordering had reversed due to an increase in the abstract wage premium. In the quantitative model of Section 4, workers with skill type $s \in (\underline{s}_{\text{old}}, \bar{s}_{\text{old}})$ face the evolution of potential wages depicted in Figure 1 (see left panel of Figure 3).

gathered experience in either of the two occupations. What is her optimal occupational choice? Assuming first that she can freely borrow against future income, the decision problem is:

$$\max_{c_T, c_{T+1}, j_T, j_{T+1}} \ln c_T + \beta \ln c_{T+1} \qquad \text{such that: } c_T + \frac{c_{T+1}}{R} = y(j_T) + \frac{y(j_T, j_{T+1})}{R}$$

where R is the exogenous interest rate, c consumption, and j_t is the occupational choice at the beginning of period t. If instead the household is exogenously prevented from borrowing, it has to hold in addition that $c_T \leq y_T$. Given that she can choose between the two occupations at the beginning of each period, four possible combinations of occupational choice exist, which in turn determine labor income:

1. $\{j_T = r, j_{T+1} = r\} \rightarrow \{y_T = \underline{h}, y_{T+1} = \bar{h}\}$
2. $\{j_T = r, j_{T+1} = a\} \rightarrow \{y_T = \underline{h}, y_{T+1} = \omega \cdot \underline{h}\}$
3. $\{j_T = a, j_{T+1} = a\} \rightarrow \{y_T = \underline{h}/\omega, y_{T+1} = \omega \cdot \bar{h}\}$
4. $\{j_T = a, j_{T+1} = r\} \rightarrow \{y_T = \underline{h}/\omega, y_{T+1} = \underline{h}\}$

Proposition 1. *If not borrowing-constrained, the household chooses the abstract occupation in $t = T$ iff*

$$R \leq \omega(\bar{h}/\underline{h} - 1) \min\left\{\frac{\bar{h}/\underline{h}}{\bar{h}/\underline{h} - 1}, \frac{\omega}{\omega - 1}\right\}.$$

In this case, she is a net borrower in $t = T$.
If borrowing-constrained, the household never chooses the abstract occupation in $t = T$.

The proof is relegated to Appendix A.1. The proposition shows that there exists a set of parameter combinations under which a household who is not borrowing-constrained optimally chooses the abstract occupation in period T, even though wages are still higher in the routine occupation. In this case, she borrows against future income in period T.[4]

The second part of Proposition 1 highlights the key inefficiency. Under no combination of parameters does the household work in the abstract occupation in T if she cannot borrow. The intuition is simple: foregoing high wages in the routine occupation today is too costly, and if wages

[4]The condition in Proposition 1 becomes more likely to hold when the interest rate R is low, such that borrowing against future income is cheap, and when the human capital spread $\bar{h}/\underline{h}$ is large, such that starting to gain experience in the abstract occupation today is valuable. ω has an ambiguous effect on the occupational choice in T. In Appendix A.1 I show that a similar proposition holds for a worker who instead of being inexperienced in period T is an experienced routine worker, i.e. $h_T = \bar{h}$.

are very high in the abstract occupation in period $T+1$, she can still switch then. That the household always chooses the routine occupation in T when borrowing constraints bind is clearly inefficient: the first result of the proposition shows that for some parameters the long-term gains from switching, in form of discounted future profits, outweigh short-term costs. In these cases, choosing the abstract occupation in T is a profitable investment.

Appendix A.1.3 extends the simple model by a stylized business cycle (time-varying aggregate productivity), while Appendix A.1.4 discusses occupation-specific returns to tenure. Appendix A.2 extends the simple model in two ways. It endogenizes wages and introduces a continuum of households who are heterogeneous with respect to their skill type. This extended set-up allows me to characterize analytically the socially optimal allocation of labor across the two occupations when the abstract wage grows over time relative to the routine wage. I show that, in line with the intuition conveyed above, in an economy in which households are hand-to-mouth, too little labor is supplied in the abstract occupation compared to the first-best. Hence, a policy that alleviates the borrowing constraint can raise output by improving the allocation of labor.

3 Empirical Evidence: Wage paths and liquid assets

This section presents empirical evidence that support and motivate the quantitative model. I review the concept of labor market polarization and the three broad occupational groups, as introduced in the earlier literature (Autor and Dorn, 2013; Jaimovich and Siu, 2020), in Appendix B.1.

3.1 Wage paths of routine workers: switchers vs. stayers

Using data from the PSID between 1976 and 2007, Cortes (2016) shows that compared to routine workers who stayed in their occupation, those who switched to manual or abstract occupations experienced faster wage growth after they had switched. In particular, to measure how wages of switchers from routine to abstract and manual occupations evolved relative to routine stayers at horizon h, he estimates the following regression:

$$\Delta_h \ln(wage_{it}) = \beta_m^h \cdot D_{imt} + \beta_a^h \cdot D_{iat} + \gamma_t^h + u_{iht}$$

Table 1: Wage changes for routine workers relative to stayers, according to direction of switch

Change in real wages between year t and year ...	To manual occupation	To abstract occupation
$t+1$	−11.2% (−8.4%)	3.4% (3.6%)
$t+2$	−14.3% (−12.6%)	5.9% (6.9%)
$t+4$	−3.5% (−3.2%)	8.5% (8.3%)
$t+10$	11.5% (13.0%)	16.3% (15.4%)

Notes: Workers who stay in routine occupations are the omitted category. Numbers show change in real wages between year t and year $t+x$, compared to workers who stayed in routine occupations reported in Cortes (2016, Table 3). Numbers in parentheses are estimates using my own sample.

where $\Delta_h \ln(wage_{it})$ is the change in the log real hourly wage of routine worker i between years t and $t+h$, D_{ijt} is a dummy variable equal to zero if the individual was working in the routine occupation in both t and $t+1$ (or t and $t+2$ for $h \geq 2$) and one if she switched from the routine occupation to occupation j, γ_t captures time fixed effects, and u_{iht} is an error term. For convenience, I reprint the baseline estimates from Cortes (2016, Table 3) of β_m^h (middle column) and β_a^h (last column) in Table 1. Numbers in parentheses show estimates from my own sample, details about which I provide below. Cortes (2016) finds that wages of workers who switched to manual occupations were 11.2% lower after one year than wages of those who stayed, but 11.5% higher after ten years. Wages of routine workers who switched to abstract occupations were on average 3.4% higher after one year compared to wages of stayers, and 16.3% higher after ten years. This pattern is consistent with a relative wage increase in manual and abstract relative to routine occupations. The quantitative model in the next section will account for these dynamics, with the average wage change after one year (first row) being a target in the calibration.

To take a closer look at the distribution of wage changes upon leaving routine occupations, I use PSID waves 1976 to 2017, following Cortes (2016) in the definition of variables and in sample selection. The only two differences of my sample compared to his are the extended time period and that, to be consistent throughout this paper, I only categorize low-skilled services occupations as manual occupations, while Cortes (2016) uses a slightly broader definition for his baseline results. I relegate all details to Appendix B.3.

I find that of all the workers who switched from routine to manual (abstract) occupations from one year to the next and for whom wages are observed in both years, 58% (42%) saw their wage decline upon switching. Hence, Cortes (2016)'s result that the one-year wage change of switchers to the abstract occupation was on average 3.4% higher than that of the stayers masks a lot of heterogeneity. While some moves from routine to abstract might have been due to career advancement, leading to wage gains, almost half of these switches have come with wage losses and are therefore less likely to represent career progression.[5]

The point here is not to claim that only leaving routine occupations can lead to (temporary) wage losses.[6] However, Cortes (2016) shows that wages of workers who switched from manual and abstract to routine occupations were either on average negative or positive but not statistically significant (the analogues of Table 1 for manual and abstract workers are re-printed in Appendix B.3.3). The takeaway here is therefore that it is predominantly for the group of routine workers that switching can be interpreted as an investment.

3.2 Hand-to-mouth shares

I have so far shown that many routine workers experienced short-term wage losses but long-term wage gains upon switching the occupation. But how many of them have faced binding liquidity constraints? To address this, I use liquid asset holdings to proxy for closeness to the borrowing constraint. Kaplan, Violante, et al. (2014) show that although only about 10% of US households have zero or close to zero net worth ("poor hand-to-mouth"), another 20% hold only very few liquid assets, while possessing some illiquid assets, e.g. a house ("wealthy hand-to-mouth"). I extend these findings, splitting the data by the three broad occupations.

I use the SCF, instead of the PSID, to shed light on hand-to-mouth shares by occupation both because the PSID has started including detailed information on asset holdings only in 1999 and because the focus of the SCF is acquiring accurate information on wealth. I use twelve waves, from 1989 to 2019. In terms of sample selection and classifying households as hand-to-mouth, I follow Kaplan, Violante, et al. (2014). In particular, I relate each household's liquid assets to current income, and classify it as hand-to-mouth if liquid asset holdings are either zero (or positive but close to zero), or if liquid assets are close to an imputed borrowing constraint, equalling one times monthly income. Appendix B.5 lays out the details.

I find that, averaged across time, 35% of households whose head was currently working in routine occupations were hand-to-mouth. Hence, for

[5] Mukoyama et al. (2023) provide empirical evidence that in the U.S. almost none of the net reallocation of labor from routine into abstract and manual occupations was driven by within-firm switches. Appendix B.3 further documents that also the subgroup of routine switchers who faced initial wage losses eventually earned higher wages than routine stayers, alleviating the concern that the results in Table 1 are driven by selection effects.

[6] In my sample, 33% of workers leaving manual and 48% of workers leaving abstract occupations saw their wage decline year-on-year.

a large fraction of routine workers binding borrowing constraints have potentially been an impediment to leaving their occupation, if such a switch entailed short-run wage losses. This might have caused them to stay in the routine occupations for a relatively long time. Appendix B.5 depicts the time series as well as confidence intervals. There, I also split up households into poor and wealthy hand-to-mouth, showing that routine workers are especially likely to be wealthy hand-to-mouth.

3.3 Liquid assets and switching behavior

In a last step, I ask whether being hand-to-mouth has had predictive power over the decision to leave the routine for the manual or abstract occupations, as the theoretical model suggests. To this end, I turn back to the PSID data, waves 1999 until 2017, because these provide information on assets (at the household level) and because, unlike the SCF, the PSID has a panel dimension and follows individuals over time. I estimate the following model:

$$\text{switch}_{i,t,t+2} = \beta \cdot HtM_{i,t} + \gamma \cdot X_{i,t} + u_{i,t} \ . \tag{1}$$

Here, $\text{switch}_{i,t,t+2}$ is a dummy variable equal to zero if worker i was employed in a routine occupation in year t and $t+2$ and equal to one if the worker was employed in a routine occupation in year t and a manual or abstract occupation in year $t+2$. I look at two-year differences as the PSID went to a bi-annual frequency in 1997. $HtM_{i,t}$ is a dummy variable indicating whether the household is hand-to-mouth in year t, and $X_{i,t}$ are control variables. $u_{i,t}$ is an exogenous error term with $\mathbb{E}[u_{i,t}] = 0$. I control for tenure in the broad occupation, a dummy indicating the region, age and its square, unionization status, married status, and a linear time trend. As in the quantitative model of the next section innate ability, or skill, will be an important determinant of switching behavior, I include in some specifications a dummy for whether individuals have received at least some years of college education.[7] Sample selection is as in Section 3.1. Standard errors are clustered at the individual level.

Table 2 shows the results. The first row of columns 1 and 2 shows the estimate of β. I find that hand-to-mouth agents are less likely to leave the routine occupation, compared to agents who hold a buffer of liquid assets. The likelihood of switching is 2.1 to 3.0 percentage points smaller for the former group than for the latter, everything else equal. This is in line with the mechanism laid out above, by which more borrowing-constrained households, for fear of temporary earnings losses, delay their move away

[7] Proxying skill by the log real hourly wage or by raw years of education yields similar results.

Table 2: Switching decision and liquid asset holdings (data)

	(1)	(2)	(3)	(4)	(5)	(6)
HtM	-0.030	-0.021	-0.032	-0.023		
	(0.012)	(0.012)	(0.012)	(0.012)		
Occupational tenure	-0.014	-0.014	-0.021	-0.020	-0.021	-0.020
	(0.00098)	(0.00095)	(0.0018)	(0.0017)	(0.0018)	(0.0017)
Skill		0.090		0.082		0.082
		(0.012)		(0.012)		(0.012)
Poor HtM					-0.047	-0.031
					(0.017)	(0.017)
Wealthy HtM					-0.024	-0.019
					(0.013)	(0.013)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4882	4882	4882	4882	4882	4882
Model	OLS	OLS	Probit	Probit	Probit	Probit

Standard errors in parentheses

Notes: Data are from the PSID, years 1999–2017. Sample selection is as in Cortes (2016), see also Appendix B.3. Dependent variable is whether or not individual leaves routine occ. between t and $t+2$. Definition of HtM status is as in Kaplan, Violante, et al. (2014). Skill is a dummy for whether the individual has received more than 12 years of education. Occ. tenure is years of uninterrupted tenure in broad occupation. Additional controls are: region dummies, age, age squared, unionization status, married status, year. Standard errors are clustered at the individual level. Columns with probit model show average marginal effect.

from the routine occupations. Also, in line with what would be expected, more years of occupational tenure make switching less likely. Skill, as measured by education, has a positive impact on switching probabilities. Columns 3 and 4 show average marginal effects in a probit model, which are similar to the OLS estimates. In columns 5 and 6 I split up the explanatory variable into poor and wealthy hand-to-mouth households. Evidently, the negative correlation is driven by both subgroups, though the relationship is a bit stronger for the poor hand-to-mouth households. I extend the empirical model in a number of ways in Appendix B.4, showing, e.g., that there is no significant relationship between liquid assets and leaving the manual or abstract occupations.

Current liquid asset holdings and future switching decisions could be jointly determined, and in fact are so in the quantitative model of Section 4. In the data, as well as in the model, there exist workers who will never choose to leave the routine occupations, given their skills. Hence, they might hold low buffers of liquid savings, as they never plan to make a career change that would temporarily depress their earnings. As a result of this simultaneity, the estimate of β in equation (1) could be biased. I therefore view these estimates only as informative correlations that support the considerations from the small model above. Also, I can use these empirical estimates as an untargeted moment to compare them to the results from a synthetic panel, simulated from the quantitative model (Appendix D). [8]

4 Quantitative Model

I now turn to the quantitative model, which contains many empirically relevant features that the analytical model abstracted from. Time t is continuous and runs forever. The economy consists of a representative firm, heterogeneous households and a government, which I now describe in detail. The exogenous driver of technological change is a falling price of ICT capital relative to the final good, $\frac{1}{q_{ict}}$. The economy is initially in its steady state in 1980 with a constant q_{ict}. It is then hit by a shock that raises q_{ict} over many years to a new, higher level, until the economy reaches a new steady state. Agents have perfect foresight over the path of q_{ict} once it is revealed in 1980. There is no aggregate risk.

4.1 Representative firm

The final good in the economy Y_t is produced according to a multiply nested CES production function (vom Lehn, 2020):

$$Y_t = \underbrace{K_{s,t}^{\alpha}}_{\text{Non-ICT capital}} \left[\mu_m N_{m,t}^{\frac{\gamma_m-1}{\gamma_m}} + (1-\mu_m)\left[\mu_a N_{a,t}^{\frac{\gamma_a-1}{\gamma_a}} + (1-\mu_a) R_t^{\frac{\gamma_a-1}{\gamma_a}}\right]^{\frac{\gamma_a(\gamma_m-1)}{(\gamma_a-1)\gamma_m}}\right]^{\frac{\gamma_m(1-\alpha)}{\gamma_m-1}} \tag{2}$$

where $R_t = [(1-\mu_r)\underbrace{K_{ict,t}^{\frac{\gamma_r-1}{\gamma_r}}}_{\text{ICT capital}} + \mu_r N_{r,t}^{\frac{\gamma_r-1}{\gamma_r}}]^{\frac{\gamma_r}{(\gamma_r-1)}}$.

Here, $N_{j,t}$ denotes effective labor employed in occupation $j \in \{m, r, a\}$. $\{\gamma_m, \gamma_r, \gamma_a, \mu_m, \mu_r, \mu_a\}$ are parameters that govern the optimal factor input shares and the elasticities of substitution. There are two types of capital, ICT and non-ICT capital (e.g. structures). The laws of motion for the two types of capital $x \in \{ict, s\}$ are, $\dot{K}_{x,t} = q_{x,t}I_{x,t} - \delta_x K_{x,t}$, where the parameters δ_x capture depreciation of the capital stock. $q_{x,t}$ denotes the amount of capital of type x that can be purchased for one unit of output at time t (Greenwood et al., 1997). The price of the final good is normalized to one. In line with Eden and Gaggl (2018), I set the relative price of non-ICT capital $\frac{1}{q_{s,t}}$ to one at all times for the remainder of this paper. All factor inputs are paid their marginal product and the firm makes zero profits.

[8] In Appendix B.4, I discuss empirical evidence in which I attempt to identify a *causal* relation between liquid assets and switching decision using data from the 2008-09 financial crisis.

4.2 Households

There exists a continuum of mass one of households who value consumption c and leisure $(1-\ell)$, with $u(c,\ell)$ denoting flow utility. Households discount the future at rate ρ and die at rate ζ, hence their effective discount rate is $\hat{\rho} \equiv \rho+\zeta$. An intensive labor supply margin is included to ensure that labor income taxation dampens labor supply. This is important when studying policies below that are financed by raising labor taxes.

To generate a share of hand-to-mouth agents as high as in the data while simultaneously modelling a production economy, I follow Kaplan, Moll, et al. (2018) in assuming that households can save in two assets: a liquid asset m and an illiquid asset $\tilde{k}$. Like them, I assume that households who die are replaced by newborns holding zero assets.[9] Households can borrow in the liquid asset up to a borrowing constraint $\underline{m}$. The liquid asset pays no interest (representing, for instance, money or low-yielding bonds), but if households borrow they have to pay an intermediation cost $\kappa > 0$. Households hold non-negative amounts of the illiquid asset, $\tilde{k} \geq 0$, which pays a return r_t. Changing $\tilde{k}$ is subject to a portfolio adjustment cost $\chi(\tilde{k}, d)$, where a positive d denotes a deposit and a negative d a withdrawal from the illiquid account.

Labor income Households choose to work in one of the three broad occupations at each instant of time t. Apart from losing occupation-specific human capital, they can costlessly switch between occupations. Denote the occupational choice by $j \in \{m, r, a\}$. Households' (pre-tax) labor income is

$$inc_j = w_j \cdot \ell \cdot y_j \ ,$$

where w_j denotes the wage per efficiency unit in occupation j, ℓ labor supply, and y_j labor productivity. The log of labor productivity is composed of:

$$\ln(y_j) = a_j \cdot (\underbrace{s}_{\text{skill}} + \underbrace{\eta}_{\text{shock}}) + \underbrace{h_j}_{\text{specif. human cap.}} + \underbrace{\epsilon}_{\text{shock}} \tag{3}$$

Each household has a fixed skill type s, which is distributed in the population according to a cumulative distribution function $F(s)$. The skill type is pre-multiplied by an occupation-specific slope parameter a_j, where I assume $0 = a_m < a_r < a_a$. This structure is borrowed from Cortes (2016), Jung and Mercenier (2014), and vom Lehn (2020) and it implies a comparative advantage of low-(high-) skilled types in the manual (abstract) occupation. This gives rise to an endogenous sorting pattern of skill types

[9] The accidental bequests of dying households are passed on to all living households in proportion to their current assets (Kaplan, Moll, et al., 2018).

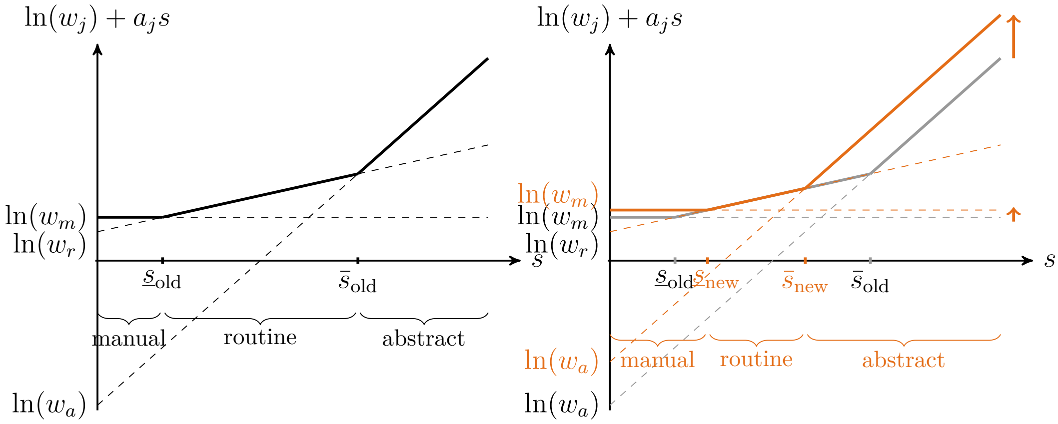

Figure 2: Wages for different skill types before (left) and after polarization (right)

into the three occupations.[10]

The left panel of Figure 2, which plots skill on the x-axis and potential earnings in each occupation on the y-axis, visualizes this. The dashed extensions of the solid lines represent hypothetical earnings of households in all three occupations, and the slopes of the lines correspond to the parameters a_j. The cut-offs $\underline{s}$ and $\bar{s}$ separate the skill space into three regions. Absent any other considerations, these cut-offs would sharply divide skill types into the three occupations in the steady state of the model. As discussed below, I introduce the shock η to capture motives for occupational switches other than technological change. On average, manual (abstract) workers earn the lowest (highest) wages in the economy, with routine workers in between.

The right panel of Figure 2 depicts skills and potential earnings after polarization has taken place, i.e. in the new steady state. The exogenous fall in the price of ICT capital leads to an endogenous increase in the abstract and manual relative to the routine wage. This shifts the skill cut-offs $\underline{s}$ and $\bar{s}$ inward, leading to a smaller set of skill types who choose the routine occupations. Hence, there are two margins of adjustment when the economy transitions to a smaller routine employment share. First, there is occupational mobility, i.e. net flows of workers switching from the routine to the manual and abstract occupation. Second, some newborn workers choose to work in the manual and abstract instead of the routine occupation, unlike the workers they have replaced. Both adjustments take place predominantly in the regions $[\underline{s}_{\text{old}}, \underline{s}_{\text{new}}]$ and $[\bar{s}_{\text{new}}, \bar{s}_{\text{old}}]$. Section 5 contains an illustrative discussion of the occupational choice of workers with skill type $\bar{s}_{\text{old}}$, when they are at the borrowing constraint and when they are away from it.

Occupation-specific human capital is captured by h_j. It can take on two

[10] Edin et al. (2023) refer to this structure as a hierarchical Roy model.

values, capturing whether the worker is inexperienced ($h_j = \underline{h}$) or experienced ($h_j = \bar{h}_j$), with $\bar{h}^j > \underline{h}\ \forall j$. For tractability, I assume that households can only be experienced ($\bar{h}_j$) in their current occupation. Once a household switches to another occupation, h is set to $\underline{h}$, implying no recall of human capital if a worker returns to her previous occupation.[11] Inexperienced households become experienced with occupation-specific Poisson intensity λ_j^h, and experienced households never become inexperienced unless they switch.

I add two shocks, η and ϵ, to the log of productivity in equation (3). η is pre-multiplied by the slope parameters a_j, and hence introduces occupational mobility for reasons other than technologocial change in a parsimonious way (Artuç et al., 2010).[12] A positive (negative) shock to η shifts households to the right (left) on the x-axis of Figure 2. A positive shock could be interpreted as a job-to-job move or a promotion, after which a household who formerly worked in the routine (manual) occupation now works in the abstract (routine) occupation. A negative shock could represent a job loss, after which a household works for a lower wage and might be unable to find another job in the abstract (routine) occupation and therefore takes a routine (manual) job. The advantage of generating steady state occupational mobility in this way is that switches up the occupational ladder are contemporaneously correlated with wage gains, and switches down the ladder usually coincide with wage losses. This allows me to target the negative (positive) average wage change of switchers from routine to manual (abstract) documented by Cortes (2016, Table 3) in the calibration (see also first row of Table 1).

The second shock, ϵ, also influences a worker's productivity, but unlike η, not her relative productivity across occupations. I include this shock to generate realistic amounts of earnings risk. The shocks evolve according to some stochastic process: $\dot{\eta}_t = \Phi_\eta(\eta_t)$, $\dot{\epsilon}_t = \Phi_\epsilon(\epsilon_t)$. Newborn households are inexperienced ($\underline{h}$) and start their lives with a draw from the invariant stationary distributions of s, η and ϵ respectively.[13]

[11] This is a common assumption (Kambourov and Manovskii, 2009a; Kikuchi and Kitao, 2020). I discuss this further in Appendix F.3.

[12] This shock resembles idiosyncratic, occupation-specific shocks that are typical in the occupational choice literature, see also Caliendo et al. (2019) and Traiberman (2019). To stay close to ealier models of labor market polarization, I choose to model the income process according to (3) instead, deviating from Cortes (2016) and vom Lehn (2020) only by introducing idiosyncratic risk and occupation-specific human capital.

[13] Several types of costs of occupational switching are not captured here for the sake of parsimony (loss of industry- or firm-specific human capital; fixed pecuniary costs, such as degrees or a licenses, moving to a new work place, buying new work clothes). These additional costs could especially discourage households close to the borrowing constraint from switching. Omitting them is therefore conservative with respect to the proposed mechanism, as they would further increase the relevance of the friction and likely make the policies in Section 6 even more powerful.

Household problem Equation (4) shows the household problem.

$$V_T(i) = \max_{\{c_t,\ell_t,d_t\},\xi} \mathbb{E}_T \int_{t=T}^{\xi} \mathrm{e}^{-\hat{\rho}(t-T)} u(c_t,\ell_t)\,\mathrm{d}t + \mathrm{e}^{-\hat{\rho}(\xi-T)} \mathbb{E}_T V^*_\xi(i) \tag{4}$$

$$\begin{aligned}\text{subj. to:}\quad & \dot{m}_t = (1-\tau_{l,t}) \cdot w_{j,t} \cdot \ell_t \cdot y_{j,t} + \mathbf{1}_{\{m_t<0\}} \cdot \kappa \cdot m_t - \chi(d_t,\tilde{k}_t) - d_t + T_t - c_t \\ & \dot{\tilde{k}}_t = r_t \tilde{k}_t + d_t \\ & m_t \geq -\underline{m},\ \tilde{k}_t \geq 0; \quad V^*_t(i) = \max_{\tilde{j}\in\{m,r,a\}} V_t(i_{-\{j,h\}}, \tilde{j}, \underline{h})\end{aligned}$$

i indicates the household state, which collects idiosyncratic state variables ($\{s,\eta,\epsilon,j,h,m,\tilde{k}\}$) and includes information about aggregate variables implied by the distribution of households over the state space. $\mathbf{1}_{\{.\}}$ denotes the indicator function. T_t is a lump-sum transfer from the government, $\tau_{l,t}$ a proportional labor income tax. $\tilde{k}_t \equiv \frac{k_{ict,t}}{q_{ict,t}} + k_{s,t}$ denotes capital in units of the final good. The interest rate r_t on $\tilde{k}$ and a no-arbitrage condition, which ensures that households are indifferent between holding either type of capital, are $r_t \equiv q_{ict,t} r_{ict,t} - (\delta_{ict} + \dot{q}_{ict,t}/q_{ict,t}) = r_{s,t} - \delta_s$. The household problem gives rise to a Hamilton-Jacobi-Bellman equation, shown in Appendix C.3.[14]

4.3 Government

The government budget constraint holds at each instant of time: $G_t + T_t = \tau_{l,t} \int_i inc_i \,\mathrm{d}\Gamma_t(i) + \dot{M}^s_t$. Here, G_t denotes government spending and $\Gamma(i)$ the cumulative distribution function of households over the idiosyncratic state space. Fiscal policy sets M^s such that the liquid asset demand from households is met, i.e. the supply of M is infinitely elastic.

4.4 Equilibrium

The equilibrium is defined in Appendix C.2.

4.5 Calibration

Relative price of ICT capital The exogenous driving force of technological change in the model is a fall in the relative price of ICT capital $\frac{1}{q_{ict}}$. To calibrate its evolution over time, I use the estimates reported in Eden and Gaggl (2018), see Figure 26 in the appendix. Agents in the model have perfect foresight over the path of q_{ict}, once it is revealed in 1980.

[14]Note also the stopping-time nature of problem (4): households choose the time $\xi \in [T,\infty)$ at which they switch the occupation. Once they do, the continuation value of the household problem is $V^*(i)$, which is equal to $V(i)$, only that the occupation j is updated and human capital reset to $\underline{h}$, while all remaining state variables ($i_{-\{j,h\}}$) stay unchanged.

Externally set parameters I calibrate a first set of parameters as in Kaplan, Moll, et al. (2018). The utility function is $u(c, \ell) = \ln(c) - \varphi \cdot \ell^{1+1/\gamma}/(1+1/\gamma)$, where γ is set to 1, and φ to 2.2. These choices ensure a Frisch elasticity of labor supply of one and an average labor supply of approximately 0.5. Households die at rate $\zeta = \frac{1}{180}$, which implies an average life span of 45 years. The unsecured borrowing limit, $\underline{m}$, is set to the average quarterly income. The portfolio adjustment cost function for the illiquid asset $\tilde{k}$ is a convex function $\chi(d, \tilde{k}) = \chi_1 \left(|d|/\tilde{k}\right)^{\chi_2} \tilde{k}$, where χ_1 and χ_2 are parameters (Alves et al., 2020). The tax rate on labor income τ_l is set to 30%, and the lump-sum transfer from the government T amounts to 6% of total output Y_t.

The share of non-ICT capital α is set to 0.34 and the depreciation rates of capital to $\delta_{ict} = 0.175$ and $\delta_s = 0.073$ annually, the average values reported in Eden and Gaggl (2018). For the parameters of the occupational production function I resort to the baseline values used by vom Lehn (2020). Normalizing $a_m = 0$, he calibrates $a_r = 0.18$ and $a_a = 0.77$. Like him, I assume that skills are standard normally distributed in the population, i.e. $s \sim N(0, 1)$. Table 12 in the appendix lists all externally calibrated parameters.

Occupation-specific human capital To calibrate the parameters of the human capital process $\{\bar{h}_m, \bar{h}_r, \bar{h}_a, \lambda_m^h, \lambda_r^h, \lambda_a^h\}$, I first estimate returns to tenure in each occupation, following Cortes (2016, Section VI C). I then choose the spreads $\bar{h}_j - \underline{h}$ and the intensities λ_h^j to minimize the mean squared deviations of the return profiles in the model from their empirical analogues. For details, see Appendix B.3.5. I find $\bar{h}_a - \underline{h} = 0.29$ and $\bar{h}_{\{m,r\}} - \underline{h} = 0.15$ as well as $\lambda_a^h = 0.023$ and $\lambda_{\{m,r\}}^h = 0.033$ to provide the best fit to the data. These estimated return profiles are broadly in line Kambourov and Manovskii (2009b), who estimate returns to occupational tenure unconditionally, i.e. across all occupations.

Productivity shock process Φ_ϵ Following Kaplan, Moll, et al. (2018), I use ϵ to target moments of earnings changes, estimated in Guvenen et al. (2021). I assume that $\Phi_\epsilon(\epsilon_t)$ is a jump-drift process $\mathrm{d}\epsilon_t = -\beta_\epsilon \epsilon_t + \mathrm{d}J_t$, with jumps arriving at rate λ_ϵ and the process drifting toward its mean of zero at rate β_ϵ. The calibration targets (Table 16 in the appendix) imply that shocks arrive on average every five years ($\lambda_\epsilon = 0.05$), with a standard deviation σ_ϵ of 1.2 and a half-life of approximately 3 years ($\beta_\epsilon = 0.05$).

Supply side I calibrate the parameters of the aggregate production function (2) as is common in the literature. I use the employment shares in the three occupational groups, as well as the share of income accruing to

Table 3: Internally calibrated parameters

Parameter	Value	Description	Target	Data (Model)
Supply side				
μ_m	0.13	PF share man.	1980 Empl. share rout.	58.5% (58.5%)
μ_r	0.94	PF share rout.	1980 Empl. share abstr.	31.6% (31.5%)
μ_a	0.69	PF share abstr.	1980 Labor share	64.0% (64.2%)
γ_m	1.67	PF elast. man.	2020 Empl. share rout.	45.0% (45.4%)
γ_r	2.47	PF elast. rout.	2020 Empl. share abstr.	42.3% (42.5%)
γ_a	0.27	PF elast. abstr.	2020 Labor share	57.0% (56.5%)
Demand side				
$\hat{\rho}$	0.018	discount rate	$\tilde{K}/Y$	2.92 (2.77)
κ	0.034	borr. wedge	M/Y	0.26 (0.30)
χ_1	0.87	portf. adj. cost	share poor HtM	0.10 (0.10)
χ_2	1.35	-	share wealthy HtM	0.20 (0.20)
Wage changes of routine switchers (1980-2020)				
$\eta_h - \eta_l$	0.78	spr. prod. grid	Avg. wage chng. switchers $r \to a$	3.4% (1.5%)
λ_η	0.02	λ prod. shock	Avg. wage chng. switchers $r \to m$	-11.2% (-10.9%)

Notes: Rates are expressed as quarterly values.

labor, both in 1980 and in 2020, to pin down the six share and elasticity parameters (μ_j, γ_j). For the 1980 employment shares I use the values reported in Autor and Dorn (2013), for the 2020 values I use own estimates from the 2019 SCF (the most recent employment shares based on CPS data reported in Kikuchi and Kitao (2020) are very similar). For the labor share I take the values reported in Eden and Gaggl (2018) for both 1980 and 2020, assuming that the labor share does not fall further between their latest observation (2013) and 2020. The calibrated values of the production function parameters are as expected (Table 3), and similar to Jaimovich, Saporta-Eksten, et al. (2021). While routine labor and ICT capital are relatively easy to substitute ($\gamma_r > 1$), abstract labor is relatively complementary to the input provided by both routine labor and ICT capital ($\gamma_a < 1$). The substitution elasticity of manual labor with the nest of abstract and routine labor is again relatively high ($\gamma_m > 1$).

Demand side As Alves et al. (2020) and Kaplan, Moll, et al. (2018), I calibrate the effective discount rate $\hat{\rho}$, the borrowing wedge κ and the parameters of the portfolio adjustment cost function $\chi(\cdot)$ targeting the ratio of liquid and illiquid assets to output in the economy, and the shares of poor and wealthy hand-to-mouth households in the initial steady state of the model. I find broadly similar values for these parameters as Kaplan, Moll, et al. (2018).

Productivity shock process Φ_η I target the initial wage change of workers who leave the routine occupations relative to those who stay (Cortes, 2016, Table 3, first column).[15] For parsimony, η follows a two-state Poisson process with symmetric transition rate λ_η between the states. This yields two parameters to target the two statistics: the spread between the states, $\eta_h - \eta_l$, and the transition rate λ_η. To obtain comparable statistics to Cortes (2016), I simulate a panel of 10,000 households along the transition path between 1980 and 2020 and perform the same regressions as him. The calibrated shock hits the workers on average every twelve years and the impact of a standard deviation shock on log productivity is significantly smaller than that of a typical shock to the other productivity variable, ϵ.

Untargeted statistics Appendix D discusses a host of untargeted statistics, e.g. a regression of switching decision on hand-to-mouth status (comparable to Section 3.3), the role of newborns in the decline of the routine employment share (Cortes, Jaimovich, Nekarda, et al., 2020) and wage changes at different horizons following occupational switches (Cortes, 2016).

5 Transition Path

I now study the transition between the steady states, demonstrating that the existence of a large share of households who hold few liquid assets causes a lag in the reallocation of labor into the abstract and manual occupation.

5.1 Individual choices

I first analyze the switching decisions of a particular set of middle-income workers whose optimal occupational choice is clearly affected by technological change. Specifically, I consider the occupational choice j_t^* of experienced routine workers ($\bar{h}_r$) of skill type $s = \bar{s}_{\text{old}}$, and ask at what point in time they decide to leave the routine for the abstract occupation.[16] Workers of this skill type predominantly work in the routine occupation before and in the abstract after polarization has shifted relative wages (see Figure 2).

The left panel of Figure 3 plots the wage that these workers can earn either as an experienced routine worker (orange) or as an inexperienced abstract worker (blue). While being an experienced routine worker clearly dominates being an inexperienced abstract worker in the 1980s, this relationship reverses in the early 2000s. The dynamic occupational choice problem therefore resembles the one in the two-period model of Section

[15] The process Φ_η could be used either to target the average wage change of switchers or the level of occupational mobility. I opt for the former, as it is central for the proposed mechanism. I report (untargeted) statistics on occupational mobility in the appendix.

[16] I condition on workers being in productivity states $\eta = \eta_h$ and $\epsilon = \mathbb{E}[\epsilon] = 0$.

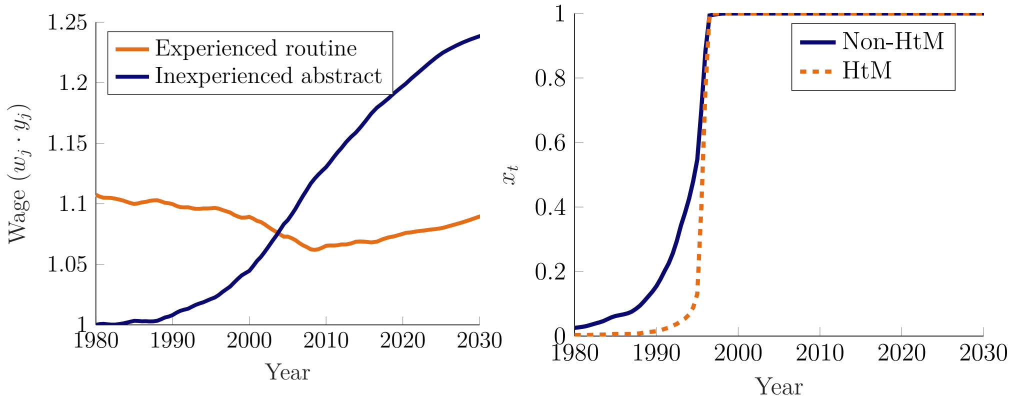

Figure 3: Experienced routine workers of skill $s = \bar{s}_{\text{old}}$: Wage (left) and likelihood of switching to the abstract occupation (right)

Notes: Left: Potential wage $w_j \cdot y_j$ for workers of skill type $s = \bar{s}_{\text{old}}$ in productivity states $\eta = \eta_h$ and $\epsilon = \mathbb{E}[\epsilon] = 0$. Orange: experienced routine workers ($w = w_{r,t}$, $h = \bar{h}_r$); blue: inexperienced abstract workers ($w = w_{a,t}$, $h = \underline{h}$). Right: Binary choice x_t of entering the abstract occupation for the same workers, averaged across wealth (m and $\tilde{k}$). Dashed orange: average over workers with $m = 0$ or $m = -\underline{m}$; solid blue: average over all remaining individuals.

2. The right panel of Figure 3 visualizes the switching choice over time, once for households that are hand-to-mouth and once for those that are not.[17] A value of $x_t = 0$ indicates that the mass of workers who would leave the routine for the abstract occupation at time t equals zero, while $x_t = 1$ implies that all of these workers would make the switch if they found themselves working in the routine occupation at time t.

As the relative wage in the abstract relative to the routine occupation rises over time, an increasing share of experienced routine workers with skill type $s = \bar{s}_{\text{old}}$ switches. However, while the share of non-hand-to-mouth workers who switch rises quickly, almost none of the hand-to-mouth households prefer a switch up until the mid-1990s. This illustrates how being close to the borrowing constraint delays investment into future earnings growth, just as it did in the tractable model of Section 2.

5.2 Aggregate implications

To assess if the distortions at the individual level have had an impact on aggregate occupational employment shares, I simulate a counterfactual transition in which all households choose their occupation like those that are rich in liquid assets and hence far away from the borrowing constraint.

Starting with the actual employment share in occupation j in the initial

[17] Formally, the right panel shows the binary choice x_t of entering the abstract (and hence leaving the routine) occupation of skill type $s = \bar{s}_{\text{old}}$, averaged across the wealth dimensions (m and $\tilde{k}$) using the stationary distribution of households in the initial steady state, i.e. $x_t = \frac{\int_{i:\mathcal{I}} \mathbf{1}_{\{j_t^* = a\}} \mathrm{d}\Gamma(i)}{\int_{i:\mathcal{I}} \mathrm{d}\Gamma(i)}$, where $i : \mathcal{I}$ indicates that individual i belongs to the set $\mathcal{I} = \{s = \bar{s}_{\text{old}}, j_{1980}^* = r, h = \bar{h}_r, \eta = \eta_h, \epsilon = 0\}$. For the orange dashed line, I replace $\Gamma(i)$ with $\Gamma(i) \cdot \mathbf{1}_{i:\{m=0 \vee m=-\underline{m}\}}$, for the blue solid line with $\Gamma(i) \cdot \mathbf{1}_{i:\{m \neq 0 \wedge m \neq -\underline{m}\}}$.

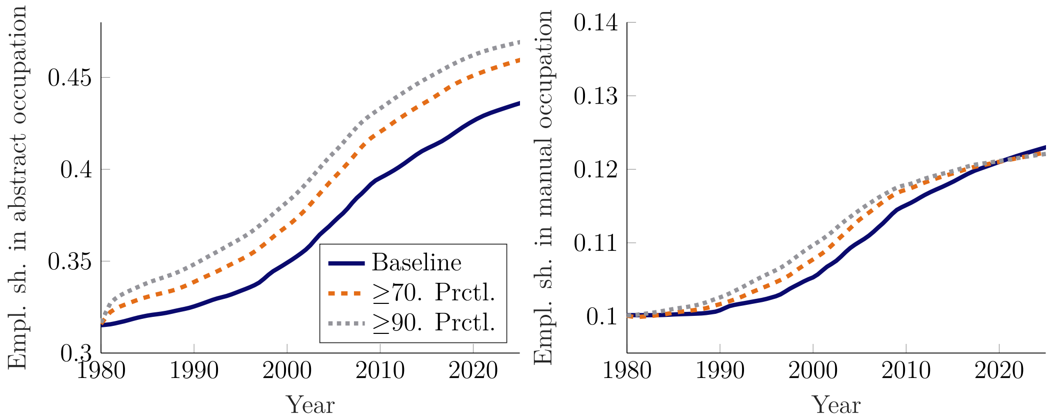

Figure 4: Baseline and counterfactual employment share in abstract (left) and manual (right) occupation

Notes: "≥70. Prctl." ("≥90. Prctl.") corresponds to counterfactual simulation in which all households choose occupations like households with liquid assets m_t greater or equal to the 70th (90th) percentile of the liquid wealth distribution.

steady state, I iterate this share forward in time as follows. For a given point in time t, I consider all households with liquid assets m_t exceeding the 70th (90th) percentile of the liquid wealth distribution. I then ask what share of these households works in occupation j, using the policy functions from the baseline transition. I use this probability to iterate forward the employment share.[18] Appendix C.5 explains the details. Note that this is not an equilibrium exercise, in contrast to the policy experiments in the next section. Here, I neither specify where the additional liquid assets come from, nor do prices respond to the changed behavior of households.

Figure 4 shows the results. Had all households acted as if they were far away from the borrowing constraint, both the manual and the abstract employment share would have been significantly higher during the transition. Between 1980 and 2020 the employment share in the abstract occupation would have been on average two (three) percentage points higher than in the baseline transition. The same effect is present for the manual employment share, though overall movements are much smaller.

6 Policy Analysis

I now study two policies, both aimed at alleviating workers' borrowing constraints, which I introduce into the model one at a time. The first is a recurrent transfer—or wage replacement—program. This is inspired by the U.S. Reemployment Trade Adjustment Assistance (RTAA) program, under which wage replacements are paid to workers who lost their job because of

[18] When solving for the transition, I discretize time into periods (Kaplan, Moll, et al., 2018). At each step, I correct for the fact that liquidity-rich households systematically differ in terms of productivity (η and ϵ) from the unconditional distribution by reweighting accordingly.

foreign trade and found reemployment at a lower wage than before. As the key friction leading to the inefficiency is the borrowing constraint, I also study a second policy, namely handing out a loan to experienced households who leave the routine occupation. The two policies differ in the incentives they provide for workers' occupational choices going forward, which will be a key issue I discuss below.

6.1 Design of the policies

All households begin their lives as non-eligible. Only experienced routine workers ($\bar{h}_r$) who switch to the abstract or manual occupation become eligible to the government program. I limit eligibility to experienced workers as the policies are supposed to target those routine workers for whom switching to the manual or abstract occupation is costly in the short run, but profitable in the long run.[19] Households are never restricted in their occupational choice after they become eligible for the program, i.e. they are allowed to move back to the routine occupation (as inexperienced workers) at a later point.

The wage replacement program entitles eligible households to a recurrent monthly transfer payment R. Workers lose eligibility and leave the program at a quarterly rate of 3%.[20] Under the loan program, eligible workers receive a one-time loan of size L from the government upon leaving the routine occupation. They then pay back the loan at a constant rate $\zeta \cdot L$ until they die. This ensures that in expectation the government breaks even on each individual loan.

There are two key differences in the design of the two programs. First, workers never have to pay back the wage replacements, unlike the loan. Second, in the spirit of RTAA, wage replacements are conditioned on working in the new (abstract or manual) occupation. If a worker decides to switch back to the routine occupation, she stays entitled but does not receive the benefit R unless she again returns to the manual or abstract occupation.

To balance its budget, the government adjusts the labor income tax $\tau_{l,t}$, while government expenditures G_t and lump-sum transfers T_t stay unchanged compared to the baseline transition. The policies start in 1980 and workers can enter the programs until 2025. All payments are stopped

[19] Moreover, for switches to the abstract occupation, at each time t I only allow skill types to become eligible among which the majority of households in t is employed in the routine occupation. This condition rules out relatively high-skilled households from being eligible, who often make immediate wage gains from switching from the routine to abstract occupation. Alternative model versions without this condition or in which I make eligibility directly conditional on experiencing a wage loss on impact deliver very similar results to the ones presented here.

[20] This matches the average rate at which inexperienced workers become experienced, and leads to an average duration in the program of about eight years.

in 2070 (an average lifetime after the last workers were allowed into the program) and then the economy slowly converges back to the same steady state as in the baseline transition without policies.

6.2 Results

I now study the introduction of the two policies in turn. I separately consider making the policy only available to switchers to the abstract and only to switchers to the manual occupation, as I find the implications of the policies to differ by the direction of switch. Since the increase in employment in the abstract occupation has been quantitatively much more important since the 1980s, I focus here on implementing the program solely for workers who switch to the abstract occupation. Targeting the policies at exits to the manual occupation is relegated to Appendix G.

6.2.1 Welfare comparison

To evaluate the policies, I measure welfare changes in consumption-equivalent units, proceeding in two steps. First, for each t, I compute the consumption-equivalent welfare change ϕ_t that makes individuals currently alive on average indifferent between living in the baseline economy and that with the policy:

$$\int_i u(\phi_t \cdot c_t^{basel}, \ell_t^{basel}) \, \mathrm{d}\Gamma_t^{basel}(i) = \int_i u(c_t^{pol}, \ell_t^{pol}) \, \mathrm{d}\Gamma_t^{pol}(i) \, . \tag{5}$$

Second, I compute the average of ϕ_t over time from $t = 1980$, when the policies are introduced, until $t = 2070$, when the last payments are stopped. For both policies, I find welfare differences to be inversely u-shaped in the size of the transfer/loan, with the optimal monthly wage replacement $R = 420\$$ and optimal loan $L = 10,000\$$. Moreover, the welfare gain under the optimal wage replacement (0.15%) is significantly larger than under the optimal loan program (0.06%).[21] Appendix E provides a graph and more details. Below I show that the average welfare gains mask large heterogeneity across birth cohorts and especially across skill types.

The policies can increase welfare for several reasons. First, they provide eligible households with insurance and improve their ability to smooth consumption. Second, at least in the case of the wage replacements, income is redistributed from tax payers to the recipients of the transfers. Since the government is utilitarian, it potentially values this redistribution. Third, by loosening the borrowing constraint for eligible households, the programs

[21]These magnitudes are similar to average welfare effects of retraining programs found in Jaimovich, Saporta-Eksten, et al. (2021) (see last line in their Table 3). See discussion below for an alternative measure of welfare.

can lead to a more efficient allocation across occupations, which can in turn cause important general equilibrium effects in the form of higher output and wages. Appendix E disentangles these different channels by 1) implementing the policies in the steady state of the model (instead of along the transition) and by 2) dissecting the total effects into partial and general equilibrium effects. Both approaches show that the general equilibrium effects generated by efficiency gains are quantitatively important.

The policies also have negative effects, which outweigh once the payments become very large. First, both policies are financed by distortionary labor taxes.[22] Second, every time an experienced routine worker switches, occupation-specific human capital gets destroyed. Hence, every additional switch induced by the policies entails a temporary resource cost. For some workers it might be efficient to incur this short-term cost, as long-run gains from switching to the abstract occupation are higher. However, some households might opt into the program only in order to pick up the payment (loan) and be able to smooth current consumption, without any intent of staying in the abstract occupation in the long run.

This helps explain why, by design, the loan program obtains smaller welfare gains than the wage replacement program. Workers are not discouraged from switching back to the routine occupation after obtaining the loan. Hence there exist many experienced routine workers who, when hit by an adverse income shock, opt into the program to collect the loan and then switch back to the routine occupation relatively quickly. In contrast, under the wage replacement program payments are stretched out over a longer time horizon and conditioned on working in the new occupation (in the spirit of RTAA). Hence, the workers who opt into this program are those who have an interest in becoming abstract workers and leaving the routine occupation for good. In short, the wage replacement program is more successful in incentivizing the right workers to leave the routine occupation than the loan program.[23]

6.2.2 General equilibrium effects under constrained optimal policies

More routine workers switching to the abstract occupation when the policies are in place puts upward (downward) pressure on the routine (abstract) wage. Moreover, Figure 5 reveals that, as some low-skilled workers switch

[22]Even in the case of the loan program, where the government breaks even in expectation on each individual loan, it has to raise taxes whenever the paybacks from earlier loan recipients are not high enough to cover the initiation of new loans. The labor tax rate needed to finance the optimal programs is shown in Figure 28.

[23]To illustrate this, Figure 27 in the appendix shows the share of program participants who have returned to the routine occupation among all currently entitled households.

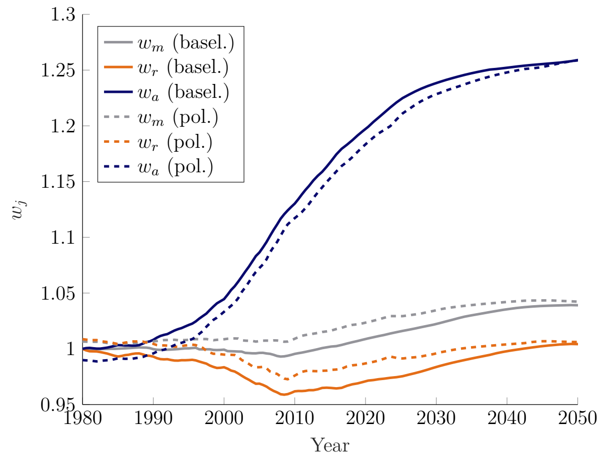

Figure 5: Wages per efficiency unit of labor

Notes: Solid lines: Baseline transition. Dashed lines: Transition with optimal wage replacement program. Wages in occupation j are relative to w_j in 1980 in the baseline.

from the manual to the routine occupation because of the now higher wages there, manual wages rise as well due to lower supply of manual labor. In sum, wages rise for a majority of the population (routine and manual workers).

Appendix E discusses in detail the effect of the policies on output, labor productivity and ICT capital. In brief, output increases on average, as the allocation of labor becomes more efficient (see also the extended tractable model in Appendix A.2). This is also reflected by an increase in labor productivity. ICT capital becomes crowded in, as less routine labor is supplied, and is up to 3% higher under the optimal wage replacement than in the baseline.

6.2.3 Welfare across cohorts and skill types

The total effect of the policies on welfare masks substantial heterogeneity both across birth cohorts and especially across skill types. To show this, Figure 6 depicts welfare changes of newborn workers by year of birth and skill type. The graphs plot the expected consumption-equivalent welfare gain of a worker of skill type s who joins the labor market in $\tilde{T} = 1980$ and 2000, and was therefore born around 1960 and 1980, respectively.[24]

[24] The expected lifetime consumption-equivalent welfare change $\bar{\phi}_{\tilde{T}}$ for a cohort entering the labor market in $\tilde{T}$ solves

$$\int_i \mathbb{E}\left[\int_{t=\tilde{T}}^{\infty} \mathrm{e}^{-\hat{\rho}(t-\tilde{T})} u\left(\bar{\phi}_{\tilde{T}} \cdot c_t^{\text{basel}}, \ell_t^{\text{basel}}\right) \mathrm{d}t\right] \mathrm{d}\bar{\Gamma}_{\tilde{T}}^{\text{basel}}(i) = \int_i \mathbb{E}\left[\int_{t=\tilde{T}}^{\infty} \mathrm{e}^{-\hat{\rho}(t-\tilde{T})} u\left(c_t^{\text{pol}}, \ell_t^{\text{pol}}\right) \mathrm{d}t\right] \mathrm{d}\bar{\Gamma}_{\tilde{T}}^{\text{pol}}(i) \tag{6}$$

where $\bar{\Gamma}_{\tilde{T}}(i)$ is the distribution of workers in $\tilde{T}$ conditional on being a newborn. Taking the average $\bar{\phi}_{\tilde{T}}$ across different cohorts $\tilde{T}$ would have been an alternative welfare criterion to select constrained optimal policies. Evaluating this is computationally very demanding, however, since the cohort is not a state variable and thus the expression inside the expectations operator is time-consuming to compute even for a single cohort. Hence, I used equation (5) as the ranking criterion above. Average gains for the cohorts

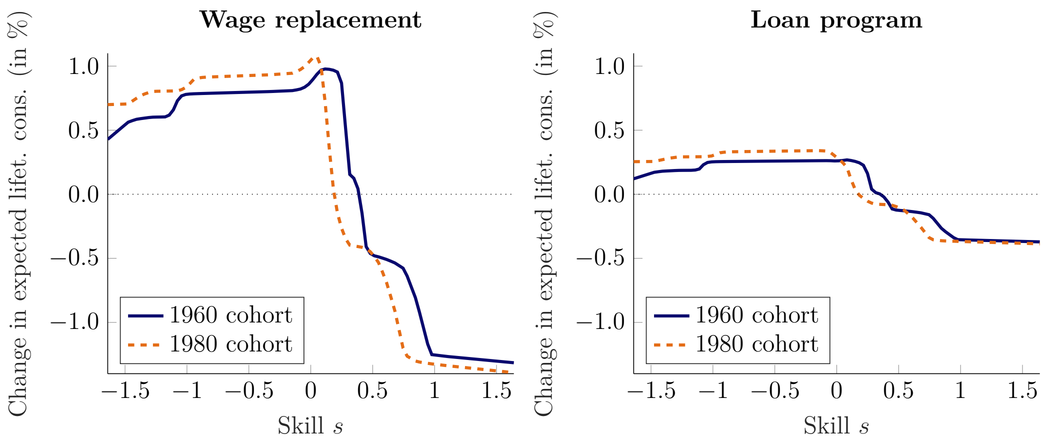

Figure 6: Expected lifetime consumption-equivalent welfare change of newborn households under constrained optimal policies compared to the baseline

Notes: $(\bar{\phi}_{\tilde{T}}(s) - 1) \cdot 100\%$ is on the y-axis, where $\bar{\phi}_{\tilde{T}}(s)$ is as defined in equation (6). Left: optimal wage replacement. Right: optimal loan program.

While the heterogeneity across birth cohorts is of minor importance, the heterogeneity across skill types is quite significant. On the one hand, low- and medium-skilled workers gain from the policies. This is both because they directly benefit from the programs by collecting the wage replacements (loans) in case they switch to the abstract occupation, and because of the rise in routine and manual wages discussed above. Quantitatively, welfare gains for medium-skilled types born in 1980 are equivalent to a 1.0% (0.3%) increase in expected lifetime consumption under the optimal wage replacement (loan) program. High-skilled workers, on the other hand, lose in terms of welfare. They usually work in the abstract occupation their entire life and hence never directly benefit from the policies. However, they have to pay higher labor taxes and are harmed by the lower abstract wage.

6.3 Robustness checks

Preference heterogeneity By assumption, all households in the model share the same preferences and are brought (close) to the borrowing constraint by idiosyncratic income shocks. This follows a large literature that models household heterogeneity in the way of Aiyagari (1994), Bewley (1983), and Huggett (1993). An alternative approach, put forward most recently by Aguiar et al. (2020), would be to assume that households are of different preference types, some more impatient than others. In Appendix F.1 I extend the model to feature a realistic degree of preference heterogeneity and show that the quantitative results are robust.

in Figure 6 (weighted by $F(s)$) are 0.15% (0.09%) for the 1960 (1980) cohort in the case of the wage replacement and 0.05% (0.06%) for the loan program.

Departure from perfect foresight That agents have perfect foresight about the path of technological change ($q_{ict} \uparrow$) is likely an exaggeration of how well households are informed about the future. However, in Appendix F.2 I show that in the opposite extreme case, in which agents are perfectly myopic, both average welfare and output gains are even larger than under perfect foresight.

Further remarks To keep the model computationally tractable, I made a few simplifying assumptions. For instance, I assumed that all human capital is lost upon an occupational switch, ruling out a later recall of it. Also, I abstracted from endogenous effort in building up human capital. I discuss these assumptions and their likely impact on the results briefly in Appendix F.3. There I also show that varying the intertemporal elasticity of substitution and the Frisch elasticity of labor supply has only small effects on the results.

7 Conclusion

The US labor market has undergone profound structural change in recent decades. In a general equilibrium model with a realistic share of hand-to-mouth households, I have demonstrated that borrowing constraints are important to understand the consequences of this labor market polarization. Welfare and output are raised by targeted government transfers or loans, and the former is preferable to the latter, as transfers better elicit those routine workers who have a long-run interest in staying in the new occupation.

In the future, more jobs will likely become automated (Edin et al., 2023), and the recent Covid-19 pandemic might even accelerate this trend (Albanesi and Kim, 2021; Chernoff and Warman, 2020). Similarly, artificial intelligence could provide a boost to automation. More generally, I have highlighted that when relative wages across occupations are on a secular trend, borrowing constraints inefficiently slow down reallocation. Such wage movements can be caused by technological change, as in this paper, or by higher demand for certain occupations due to shifting preferences. Higher demand for jobs related to the green transition, which Saussay et al. (2022) show to require a specific skill set, is one example. This article provides a rationale for supporting the needed labor reallocation with targeted government programs.

References

Acemoglu, Daron and David H Autor (2011). *Skills, tasks and technologies: Implications for employment and earnings*. Vol. 4. PART B. Elsevier B.V., pp. 1043–1171.

Acemoglu, Daron and Pascual Restrepo (2018). “The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment”. In: *American Economic Review* 108.6, pp. 1488–1542.

Achdou, Yves, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lionse, and Benjamin Moll (2022). “Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach [On the Existence and Uniqueness of Stationary Equilibrium in Bewley Economies with Production]”. In: *Review of Economic Studies* 89.1, pp. 45–86.

Aguiar, Mark A., Mark Bils, and Corina Boar (2020). *Who Are the Hand-to-Mouth?* NBER Working Papers 26643. National Bureau of Economic Research, Inc.

Aiyagari, S. Rao (1994). “Uninsured Idiosyncratic Risk and Aggregate Saving”. In: *The Quarterly Journal of Economics* 109.3, pp. 659–684.

Albanesi, Stefania and Jiyeon Kim (2021). “Effects of the COVID-19 Recession on the US Labor Market: Occupation, Family, and Gender”. In: *Journal of Economic Perspectives* 35.3, pp. 3–24.

Alves, Felipe, Greg Kaplan, Benjamin Moll, and Giovanni L. Violante (2020). “A Further Look at the Propagation of Monetary Policy Shocks in HANK”. In: *Journal of Money, Credit and Banking* 52.S2, pp. 521–559.

Artuç, Erhan, Shubham Chaudhuri, and John McLaren (2010). “Trade Shocks and Labor Adjustment: A Structural Empirical Approach”. In: *American Economic Review* 100.3, pp. 1008–1045.

Autor, David H and David Dorn (2013). “The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market”. In: *American Economic Review* 103.5, pp. 1553–1597.

Autor, David H, Frank Levy, and Richard J Murnane (2003). “The Skill Content of Recent Technological Change: An Empirical Exploration”. In: *The Quarterly Journal of Economics* November, pp. 1279–1333.

Beraja, Martin and Nathan Zorzi (2022). *Inefficient Automation*. NBER Working Papers 30154. National Bureau of Economic Research, Inc.

Bernhardt, Dan and David Backus (1990). “Borrowing Constraints, Occupational Choice, and Labor Supply”. In: *Journal of Labor Economics* 8.1, pp. 145–173.

Bewley, Truman (1983). “A Difficulty with the Optimum Quantity of Money”. In: *Econometrica* 51.5, pp. 1485–1504.

Caliendo, Lorenzo, Maximiliano Dvorkin, and Fernando Parro (2019). “Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock”. In: *Econometrica* 87.3, pp. 741–835.

Chernoff, Alex W. and Casey Warman (2020). *COVID-19 and Implications for Automation*. NBER Working Papers 27249. National Bureau of Economic Research, Inc.

Cloyne, James, Clodomiro Ferreira, and Paolo Surico (2020). “Monetary Policy when Households have Debt: New Evidence on the Transmission Mechanism”. In: *Review of Economic Studies* 87.1, pp. 102–129.

Cortes, Guido Matias (2016). "Where Have the Middle-Wage Workers Gone? A Study of Polarization Using Panel Data". In: *Journal of Labor Economics* 34.1.
Cortes, Guido Matias and Giovanni Gallipoli (2018). "The Costs of Occupational Mobility: An Aggregate Analysis". In: *Journal of the European Economic Association* 16.2, pp. 275–315.
Cortes, Guido Matias, Nir Jaimovich, Christopher J. Nekarda, and Henry E. Siu (2020). "The dynamics of disappearing routine jobs: A flows approach". In: *Labour Economics* 65, p. 101823.
Cortes, Guido Matias, Nir Jaimovich, and Henry E. Siu (2017). "Disappearing routine jobs: Who, how, and why?" In: *Journal of Monetary Economics* 91, pp. 69–87.
Dorn, David (2009). "Essays on Inequality, Spatial Interaction, and the Demand for Skills". Doctoral dissertation. University of St. Gallen.
Eden, Maya and Paul Gaggl (2018). "On the Welfare Implications of Automation". In: *Review of Economic Dynamics* 29, pp. 15–43.
Edin, Per-Anders, Tiernan Evans, Georg Graetz, Sofia Hernnäs, and Guy Michaels (2023). "Individual Consequences of Occupational Decline". In: *The Economic Journal* 133.654, pp. 2178–2209.
Evans, David S and Boyan Jovanovic (1989). "An Estimated Model of Entrepreneurial Choice under Liquidity Constraints". In: *Journal of Political Economy* 97.4, pp. 808–827.
Galor, Oded and Joseph Zeira (1993). "Income Distribution and Macroeconomics". In: *Review of Economic Studies* 60.1, pp. 35–52.
Gathmann, Christina and Uta Schönberg (2010). "How General Is Human Capital? A Task-Based Approach". In: *Journal of Labor Economics* 28.1.
Goos, Maarten, Alan Manning, and Anna Salomons (2014). "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring". In: *American Economic Review* 104.8, pp. 2509–2526.
Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell (1997). "Long-Run Implications of Investment-Specific Technological Change". In: *American Economic Review* 87.3, pp. 342–362.
Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song (2021). "What Do Data on Millions of U.S. Workers Reveal About Lifecycle Earnings Dynamics?" In: *Econometrica* 89.5, pp. 2303–2339.
Hershbein, Brad and Lisa B. Kahn (2018). "Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings". In: *American Economic Review* 108.7, pp. 1737–1772.
Huggett, Mark (1993). "The risk-free rate in heterogeneous-agent incomplete-insurance economies". In: *Journal of Economic Dynamics and Control* 17.5-6, pp. 953–969.
Jaimovich, Nir, Itay Saporta-Eksten, Henry Siu, and Yaniv Yedid-Levi (2021). "The macroeconomics of automation: Data, theory, and policy analysis". In: *Journal of Monetary Economics* 122, pp. 1–16.
Jaimovich, Nir and Henry E. Siu (2020). "Job Polarization and Jobless Recoveries". In: *The Review of Economics and Statistics* 102.1, pp. 129–147.
Jung, Jaewon and Jean Mercenier (2014). "Routinization-Biased Technical Change And Globalization: Understanding Labor Market Polarization". In: *Economic Inquiry* 52.4, pp. 1446–1465.

Kambourov, Gueorgui and Iourii Manovskii (2008). “Rising occupational and industry mobility in the United States: 1968-97”. In: *International Economic Review* 49.1, pp. 41–79.
— (2009a). “Occupational Mobility and Wage Inequality”. In: *Review of Economic Studies* 76.2, pp. 731–759.
— (2009b). “Occupational specificity of human capital”. In: *International Economic Review* 50.1, pp. 63–115.
Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante (2018). “Monetary Policy According to HANK”. In: *American Economic Review* 108.3, pp. 697–743.
Kaplan, Greg, Giovanni L. Violante, and Justin Weidner (2014). “The Wealthy Hand-to-Mouth”. In: *Brookings papers on Economic Activity* Spring, pp. 77–153.
Kikuchi, Shinnosuke and Sagiri Kitao (2020). *Welfare Effects of Polarization: Occupational Mobility over the Life-cycle*. Tech. rep.
Lochner, Lance and Alexander Monge-Naranjo (2012). “Credit Constraints in Education”. In: *Annual Review of Economics* 4.1, pp. 225–256.
— (2016). “Student Loans and Repayment: Theory, Evidence, and Policy”. In: vol. 5. Handbook of the Economics of Education. Elsevier. Chap. 8, pp. 397–478.
Mian, Atif, Kamalesh Rao, and Amir Sufi (2013). “Household Balance Sheets, Consumption, and the Economic Slump”. In: *The Quarterly Journal of Economics* 128.4, pp. 1687–1726.
Mian, Atif and Amir Sufi (2011). “House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis”. In: *American Economic Review* 101.5, pp. 2132–2156.
Moll, Benjamin, Lukasz Rachel, and Pascual Restrepo (2022). “Uneven Growth: Automation’s Impact on Income and Wealth Inequality”. In: *Econometrica* 90.6, pp. 2645–2683.
Mukoyama, Toshi, Satoshi Tanaka, and Naoki Takayama (2023). “Occupational Reallocations within and across Firms: Implications for Labor Market Polarization”.
Saussay, Aurelien, Misato Sato, Francesco Vona, and Layla O’Kane (2022). *Who’s fit for the low-carbon transition? Emerging skills and wage gaps in job ad data*. LSE Research Online Documents on Economics 117254. London School of Economics and Political Science, LSE Library.
Sullivan, Paul (2010). “Empirical evidence on occupation and industry specific human capital”. In: *Labour Economics* 17.3, pp. 567–580.
Traiberman, Sharon (2019). “Occupations and Import Competition: Evidence from Denmark”. In: *American Economic Review* 109.12, pp. 4260–4301.
vom Lehn, Christian (2020). “Labor Market Polarization, the Decline of Routine Work, and Technological Change: A Quantitative Analysis”. In: *Journal of Monetary Economics* 110, pp. 62–80.

A Details and Extensions of Tractable Model

A.1 Proofs and time-varying productivity

A.1.1 Proofs for Simple Model

Note that the model of Section 2 can be more generally described as a model where the household arrives in T either with some prior occupation $j_{T-1} \in \{r, a\}$ or as a newborn with no prior occupation, $j_{T-1} = \emptyset$. The prior occupation is then the only relevant state variable of the household. If she works in the same occupation for two consecutive periods, she is an experienced household ($\bar{h}$). If she switches the occupation or if she has no prior occupation (e.g. because she newly arrived to the labor market) she is inexperienced ($\underline{h}$). Therefore, $h_t = \bar{h} \cdot \mathbf{1}_{\{j_t = j_{t-1}\}} + \underline{h} \cdot \mathbf{1}_{\{j_t \neq j_{t-1}\}}$.

Experienced routine worker Let me first state the proposition regarding the occupational choice of a worker who was priorly employed in the routine occupation, i.e. with $j_{T-1} = r$. The potential occupational choices are:

1. $\{j_T = r, j_{T+1} = r\} \rightarrow \{y_T = \bar{h}, y_{T+1} = \bar{h}\}$
2. $\{j_T = r, j_{T+1} = a\} \rightarrow \{y_T = \bar{h}, y_{T+1} = \omega \cdot \underline{h}\}$
3. $\{j_T = a, j_{T+1} = a\} \rightarrow \{y_T = \underline{h}/\omega, y_{T+1} = \omega \cdot \bar{h}\}$
4. $\{j_T = a, j_{T+1} = r\} \rightarrow \{y_T = \underline{h}/\omega, y_{T+1} = \underline{h}\}$

Proposition 2. *If not borrowing-constrained, the household switches to the abstract occupation in $t = T$ iff*

$$R \leq \omega - \frac{\max\left\{\omega, \bar{h}/\underline{h}\right\} - 1}{\bar{h}/\underline{h} - 1/\omega}$$

In this case, she is a net borrower in $t = T$.
If borrowing-constrained, the household never switches to the abstract occupation in $t = T$.

The result resembles the one for the inexperienced worker in Section 2. This was to be expected: the only difference in the case of the experienced routine worker is that choosing the abstract occupation in period T is even more costly because of the already accumulated occupation-specific human capital. This also explains why a growing human capital spread $\bar{h}/\underline{h}$ does not unambiguously make the condition in Proposition 2 more likely to hold (like it did in Proposition 1), as it further raises the cost of leaving the routine occupation. As was the case for inexperienced workers, experienced workers who are borrowing-constrained never choose the abstract occupation in T.

A.1.2 Proofs of Propositions 1 and 2

Unconstrained households I first prove the result regarding the unconstrained household of Proposition 1 in the main text.
For the household to choose the abstract occupation in period T, it must hold that option 3 yields higher lifetime income than the maximum of options 1 and 2, i.e.

$$\underline{h} + \frac{1}{R}\max\{\bar{h}, \omega\underline{h}\} \leq \underline{h}/\omega + \frac{\omega}{R}\bar{h}$$

Case 1: $\bar{h}/\underline{h} \geq \omega$

$$\begin{aligned} \underline{h} + \frac{1}{R}\bar{h} &\leq \underline{h}/\omega + \frac{\omega}{R}\bar{h} \\ R\left(1 - \frac{1}{\omega}\right) &\leq (\omega - 1)\bar{h}/\underline{h} \\ R &\leq \frac{\omega - 1}{1 - \frac{1}{\omega}}\bar{h}/\underline{h} = \omega\bar{h}/\underline{h} = \omega(\bar{h}/\underline{h} - 1)\frac{\bar{h}/\underline{h}}{\bar{h}/\underline{h} - 1} \end{aligned}$$

Case 2: $\bar{h}/\underline{h} < \omega$

$$\begin{aligned} \underline{h} + \frac{\omega}{R}\underline{h} &\leq \underline{h}/\omega + \frac{\omega}{R}\bar{h} \\ R\left(1 - \frac{1}{\omega}\right) &\leq \omega(\bar{h}/\underline{h} - 1) \\ R &\leq \omega\frac{\bar{h}/\underline{h} - 1}{1 - \frac{1}{\omega}} = \omega(\bar{h}/\underline{h} - 1)\frac{\omega}{\omega - 1} \end{aligned}$$

Rearranging yields the desired result, which I reprint here for convenience:

$$R \leq \omega(\bar{h}/\underline{h} - 1)\min\left\{\frac{\bar{h}/\underline{h}}{\bar{h}/\underline{h} - 1}, \frac{\omega}{\omega - 1}\right\}.$$

The proof of Proposition 2 regarding the unconstrained households follows analogously.

It is also straightforward to verify that whenever the unconstrained household chooses option 3, she borrows against future income. Too see this, note that consumption in period T is $c_T = \frac{1}{1+\beta}\left(\underline{h}/\omega + \frac{\omega}{R}\bar{h}\right)$, and hence savings in period T are $s = \frac{1}{1+\beta}\left(\beta\underline{h}/\omega - \frac{\omega}{R}\bar{h}\right)$. Assume for a contradiction that savings in T are positive

Case 1: $\bar{h}/\underline{h} \geq \omega$

$$\begin{aligned} & \beta\underline{h}/\omega > \frac{\omega}{R}\bar{h} \\ \Rightarrow \quad & \underline{h}/\omega > \frac{\omega}{R}\bar{h} \\ \Rightarrow \quad & R > \omega^2\bar{h}/\underline{h} > \omega\bar{h}/\underline{h} \quad \lightning \end{aligned}$$

Case 2: $\bar{h}/\underline{h} < \omega$

$$\begin{aligned} & \beta\underline{h}/\omega > \frac{\omega}{R}\bar{h} \\ \Rightarrow \quad & \underline{h}/\omega > \frac{\omega}{R}\bar{h} \\ \Rightarrow \quad & R > \omega^2\bar{h}/\underline{h} > \omega^2\frac{\bar{h}/\underline{h} - 1}{\omega - 1} = \omega\frac{\bar{h}/\underline{h} - 1}{1 - \frac{1}{\omega}} \quad \lightning \end{aligned}$$

Both contradictions arise from the conditions derived in the first part of the proof.

Constrained households Next I prove that if the household is prevented from borrowing, she never chooses option 3.

Suppose first that in the optimal solution the borrowing constraint is not binding. In this case, the problem of the constrained household is identical to that of the unconstrained household, whose solution I have just derived. We know, however, that if parameters are such that option 3 (choosing abstract in $t = T$) is optimal she is a net borrower, a contradiction to the assumption that the constraint is not binding.

Now suppose the borrowing constraint is binding and the constrained household consumes her current income in both periods. For her to find it optimal to choose option 3 it must hold that

$$\ln(\underline{h}) + \beta \max\left\{\ln\left(\bar{h}\right), \ln(\omega\underline{h})\right\} \leq \ln(\underline{h}/\omega) + \beta \ln\left(\omega\bar{h}\right)$$

Case 1: $\bar{h}/\underline{h} \geq \omega$

$$\begin{aligned} \ln(\underline{h}) + \beta\ln\left(\bar{h}\right) &\leq \ln(\underline{h}) - \ln(\omega) + \beta\ln(\omega) + \beta\ln\left(\bar{h}\right) \\ 0 &\leq \beta - 1 \quad \lightning \end{aligned}$$

The contradiction follows from the fact that I have assumed (strict) discounting, i.e. $\beta < 1$.

Case 2: $\bar{h}/\underline{h} < \omega$

$$\begin{aligned} \ln(\underline{h}) + \beta\ln(\omega) + \beta\ln(\underline{h}) &\leq \ln(\underline{h}) - \ln(\omega) + \beta\ln(\omega) + \beta\ln\left(\bar{h}\right) \\ \ln(\omega) &\leq \beta\ln\left(\bar{h}/\underline{h}\right) \\ \omega &\leq (\bar{h}/\underline{h})^{\beta} \quad \lightning \end{aligned}$$

The proof of Proposition 2 regarding the constrained households follows analogously.

A.1.3 Time-varying aggregate productivity

Consider now the introduction of an additional variable $Z > 1$ that augments all wages in period $T+1$, i.e.

$$\begin{aligned} w_{r,T} &= 1, & w_{r,T+1} &= Z \\ w_{a,T} &= \frac{1}{\omega}, & w_{a,T+1} &= Z\omega \end{aligned}$$

The following proposition holds regarding the occupational choice of an inexperienced household, and can be proved analogously to Proposition 1 above.

Proposition 3. *If not borrowing-constrained, the household chooses the abstract occupation in $t = T$ iff*

$$\frac{R}{Z} \leq \omega(\bar{h}/\underline{h} - 1)\min\left\{\frac{\bar{h}/\underline{h}}{\bar{h}/\underline{h} - 1}, \frac{\omega}{\omega - 1}\right\}.$$

In this case, she is a net borrower in $t = T$.
If borrowing-constrained, the household never chooses the abstract occupation in $t = T$.

This result is identical to Proposition 1, except for the fact that the gross interest rate R is dampened by the factor Z. Put differently, a higher Z makes choosing the abstract occupation today more likely. Intuitively, if wages today are relatively low compared to tomorrow, the opportunity cost of choosing the abstract occupation today is also relatively low. This is in line with empirical evidence on the business cycle patters of the aggregate decline in routine labor (Hershbein and Kahn, 2018; Jaimovich and Siu, 2020).[25]

[25] The cyclical patterns of the hand-to-mouth share have so far, to the best of my knowledge, not been studied in detail. Note, however, that the share of hand-to-mouth households in the US economy depicted in Figure 10 or in Kaplan, Violante, et al. (2014) does not display any notable pro- or countercyclical pattern.

A.1.4 Occupation-specific returns to tenure

Even if on aggregate relative wages between occupations were constant, wage paths across occupations could still be differently steep from the viewpoint of the individual if experience is more highly rewarded in one occupation than in the other. For instance, returns to tenure might be higher in the abstract than in the routine occupation, in which case the ratio $\bar{h}_j/\underline{h}$ would be occupation-specific, with $\bar{h}_a/\underline{h} > \bar{h}_r/\underline{h}$ (e.g. a manager's occupational experience being more highly rewarded than that of a bookkeeper). Then, even if wages per efficiency unit in the two occupations were constant over time, from the individual viewpoint income paths would resemble those in Figure 1, and borrowing-constrained households might end up inefficiently choosing the routine occupation in $t = T$. While I allow for heterogeneous tenure profiles in the full model in Section 4, the focus of this paper is technological change and the time-varying wages it induces across occupations.[26]

A.2 Tractable model extended by endogenous wages and skill distribution

In the following I extend the simple model of Section 2 in two ways. First, I endogenize wages. Second, I allow for a continuum of households with heterogeneous skill type. I first use the extended environment to characterize the first-best, i.e the efficient allocation of labor across occupations when there is relative wage growth in the abstract relative to the routine occupation. I then study the allocation in the decentralized equilibrium, once when households can freely borrow against future income and once when everyone is hand-to-mouth. In the former case, the allocation of labor coincides with that chosen by the planner. In the latter, relative to the first-best, too few households work in the abstract and too many in the routine occupation.

A.2.1 Environment

Time t is discrete and runs until period $T+1$. There exists a representative firm that uses two inputs, routine labor N_r and abstract labor N_a. It

[26] Bernhardt and Backus (1990) show that borrowing constraints in combination with differential earnings slopes across occupations lead to inefficient job choices. For a treatment of how returns to college and its cost shape the optimal design of student loan programs in the presence of borrowing constraints, see Lochner and Monge-Naranjo (2012, 2016). In Appendix E I show that the policies raise welfare significantly more along the transition path, which features time-varying relative wages between occupations, than they do in the steady state, in which occupation-specific returns to tenure are present, but relative wages across occupations are constant.

produces output according to the production function $F(N_r, N_a; q)$ where q is the exogenous level of technology.

There exists a continuum of measure one of households. Households derive utility from consumption $u(c)$ and provide one unit of labor inelastically in one of the two occupations. A household is characterized by her fixed skill type s and her previous period's occupation j_{t-1}. Households face a constant probability of dying $\pi \in (0, 1)$. A dead household is replaced by a newborn who has no previous occupation, $j_{t-1} = \emptyset$. All households die with certainty after period $T + 1$. Conditional on survival, households discount future utility by a factor $\beta \in (0, \frac{1}{1-\pi})$. Denote the effective discount rate of households $\hat{\beta} \equiv \beta \cdot (1 - \pi)$, and hence $\hat{\beta} \in (0, 1)$.

Skill is distributed in the population according to some probability distribution function $s \sim g(s)$. Households accumulate occupation-specific human capital $h \in \{\underline{h}, \bar{h}\}$, with $0 < \underline{h} < \bar{h}$. If a household stays in the same occupation in two consecutive periods, i.e. $j_t = j_{t-1}$ she is an experienced worker in period t ($h_t = \bar{h}$). If she switches occupations or if she is a newborn with no prior occupation, she is inexperienced ($h_t = \underline{h}$), as $j_t \neq j_{t-1}$. The functions $\phi_j(s)$ with $j \in \{r, a\}$ map skill into productivity in the respective occupation. Hence if a household works in occupation j at time t, her labor supply is $\phi_j(s) \cdot h_t$. I make the following assumptions:

Assumptions 1.

- *$F(\cdot)$ is continuously differentiable, with $F_{N_j} > 0$ for $j = r, a$*
- *Without loss of generality, an increase in q raises the marginal product of abstract labor relative to that of routine labor, $\frac{\partial \frac{F_{N_a}}{F_{N_r}}}{\partial q} > 0$*
- *$u(\cdot)$ is twice continuously differentiable, with $u_c > 0$, $u_{cc} < 0$*
- *$g(s)$ has positive mass on the whole real line: $\forall s \in \mathbb{R}, g(s) > 0$*
- *$\phi(s) \equiv \frac{\phi_a(s)}{\phi_r(s)}$ is continuously differentiable, with $\phi'(s) > 0$, $\lim_{s \to -\infty} = 0$ and $\lim_{s \to \infty} = \infty$*

The last assumption implies that high-skilled types have a comparative advantage in the abstract occupation, and low-skilled types in the routine occupation.[27]

I assume that the economy can transfer resources on aggregate to the next period by lending and borrowing from the rest of the world at rate $R \leq 1/\beta$. I assume that holdings of the foreign asset are zero when entering time T and all debt taken on in T has to be repaid in period $T + 1$.

[27]This assumption is common in the literature (Cortes, 2016; vom Lehn, 2020) and it is met in the full model of the main text, where I calibrate $\phi_j(s) = \exp(a_j s)$ with $0 = a_m < a_r < a_a$.

A.2.2 Social planner

I begin by studying the first-best allocation of labor that arises as the solution to the social planner's problem.

No technological change Assume first that there is no technological change: $\forall t, q_t = q$. In period T, the planner needs to allocate consumption and an occupation to each household for the remaining two periods T and $T+1$. Note first that since I have assumed concave utility, the planner assigns the same level of consumption C_t to each household. Second, the assumption on the productivity functions $\phi_j(s)$ and the fact that marginal products of labor are constant imply that there is no occupational mobility and that there exists a cut-off S that separates the skill space into low-skilled (who work in the routine occupation their entire life) and high-skilled (abstract occupation).

$$\max_{S, C_T, C_{T+1}} u(C_T) + \beta \cdot u(C_{T+1})$$

$$\text{subject to:}$$

$$C_T + \frac{C_{T+1}}{R} = Y_T + \frac{Y_{T+1}}{R}$$

$$\text{and for } t = T, T+1:$$

$$Y_t = Y = F(N_r, N_a; q)$$

$$N_r = \int_{-\infty}^{S} \left[\pi \cdot (\phi_r(s) \cdot \underline{h}) + (1-\pi) \cdot (\phi_r(s) \cdot \bar{h})\right] g(s) \mathrm{d}s$$

$$N_a = \int_{S}^{\infty} \left[\pi \cdot (\phi_a(s) \cdot \underline{h}) + (1-\pi) \cdot (\phi_a(s) \cdot \bar{h})\right] g(s) \mathrm{d}s$$

The optimal skill cut-off $\mathbf{S}$, which maximizes output, is characterized by:

$$F_{N_r} \cdot \phi_r(\mathbf{S}) = F_{N_a} \cdot \phi_a(\mathbf{S})$$

$$\Leftrightarrow F_{N_r} = F_{N_a} \cdot \phi(\mathbf{S}) \tag{7}$$

Intuitively, this condition ensures that the worker with skill type $s = \mathbf{S}$ is equally productive in both occupations.

Technological change Now assume that technology is constant up until time T. In period T, however, it becomes known that $q_{T+1} > q_T$ and hence the marginal product of abstract labor relative to that of routine labor grows. It is thus optimal to reallocate workers to the abstract occupation, i.e. even those with skill below the optimal skill cut-off absent technological change $\mathbf{S}$. Moreover, it might be optimal to start doing so already in period T such that workers can start building up occupation-specific human capital.

The planner now needs to choose four cut-off levels, one for each combination of human capital and time period. For $t = T, T+1$, she needs to decide up to which skill level $S_{1,t}$ inexperienced workers work in the routine occupation. She also needs to decide up to which level $S_{2,t}$ experienced routine workers keep working in the routine occupation. Note that in any optimal allocation it must hold that $S_{1,t} \leq S_{2,t}$. In words, young workers are at least as likely to enter the abstract occupation as experienced routine workers, since the latter have already accumulated occupation-specific human capital.[28]

$$\max_{\{S_{1,t}, S_{2,t}, C_t\}_{t=T,T+1}} \quad u(C_T) + \beta \cdot u(C_{T+1})$$

$$\text{subject to:}$$

$$C_T + \frac{C_{T+1}}{R} = Y_T + \frac{Y_{T+1}}{R}$$

$$Y_t = F(N_{r,t}, N_{a,t}; q_t), \text{ for } t = T, T+1$$

$$N_{r,T} = \pi \int_{-\infty}^{S_{1,T}} (\phi_r(s)\underline{h}) g(s) \mathrm{d}s + (1-\pi) \int_{-\infty}^{S_{2,T}} (\phi_r(s)\bar{h}) g(s) \mathrm{d}s$$

$$N_{a,T} = \pi \int_{S_{1,T}}^{\infty} (\phi_a(s)\underline{h}) g(s) \mathrm{d}s + (1-\pi) \left[\int_{\mathbf{S}}^{\infty} (\phi_a(s)\bar{h}) g(s) \mathrm{d}s + \int_{S_{2,T}}^{\mathbf{S}} (\phi_a(s)\underline{h}) g(s) \mathrm{d}s \right]$$

and equations for $N_{r,T+1}$ $N_{a,T+1}$, for which one needs to differentiate three cases, depending on the ordering of $S_{1,T}$, $S_{2,T}$ and $S_{2,T+1}$:[29]

Case 1:

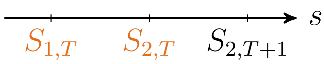

[28]This is easy to prove by contradiction. Assume that $S_{1,t} > S_{2,t}$. But then one can move the cut-offs closer together, achieving the same amount of abstract labor supply but more routine labor supply. This will become evident formally when I derive the optimal cut-offs below.

[29]The position of $S_{1,T+1}$ relative to $S_{1,T}$ and $S_{2,T}$ is inconsequential in what follows.

$$N_{r,T+1} = \pi \int_{-\infty}^{S_{1,T+1}} (\phi_r(s)\underline{h})g(s)\mathrm{d}s + (1-\pi)\left[(1-\pi)\int_{-\infty}^{S_{2,T}} (\phi_r(s)\bar{h})g(s)\mathrm{d}s + \pi \int_{-\infty}^{S_{1,T}} (\phi_r(s)\bar{h})g(s)\mathrm{d}s\right]$$

$$N_{a,T+1} = \pi \int_{S_{1,T+1}}^{\infty} (\phi_a(s)\underline{h})g(s)\mathrm{d}s + (1-\pi)\left[(1-\pi)\int_{S_{2,T}}^{\infty} (\phi_a(s)\bar{h})g(s)\mathrm{d}s + \pi \int_{S_{1,T}}^{\infty} (\phi_a(s)\bar{h})g(s)\mathrm{d}s\right]$$

Case 2:

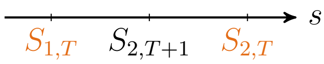

$$N_{r,T+1} = \pi \int_{-\infty}^{S_{1,T+1}} (\phi_r(s)\underline{h})g(s)\mathrm{d}s + (1-\pi)\left[(1-\pi)\int_{-\infty}^{S_{2,T+1}} (\phi_r(s)\bar{h})g(s)\mathrm{d}s + \pi \int_{-\infty}^{S_{1,T}} (\phi_r(s)\bar{h})g(s)\mathrm{d}s\right]$$

$$\begin{aligned} N_{a,T+1} = \pi \int_{S_{1,T+1}}^{\infty} (\phi_a(s)\underline{h})g(s)\mathrm{d}s + (1-\pi)\Bigg[&(1-\pi)\int_{S_{2,T+1}}^{S_{2,T}} (\phi_a(s)\underline{h})g(s)\mathrm{d}s \\ &+ (1-\pi)\int_{S_{2,T}}^{\infty} (\phi_a(s)\bar{h})g(s)\mathrm{d}s + \pi \int_{S_{1,T}}^{\infty} (\phi_a(s)\bar{h})g(s)\mathrm{d}s\Bigg] \end{aligned}$$

Case 3:

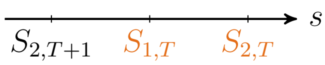

$$N_{r,T+1} = \pi \int_{-\infty}^{S_{1,T+1}} (\phi_r(s)\underline{h})g(s)\mathrm{d}s + (1-\pi)\int_{-\infty}^{S_{2,T+1}} (\phi_r(s)\bar{h})g(s)\mathrm{d}s$$

$$\begin{aligned} N_{a,T+1} = \pi \int_{S_{1,T+1}}^{\infty} (\phi_a(s)\underline{h})g(s)\mathrm{d}s + (1-\pi)\Bigg[&\pi\int_{S_{2,T+1}}^{S_{1,T}} (\phi_a(s)\underline{h})g(s)\mathrm{d}s + (1-\pi)\int_{S_{2,T+1}}^{S_{2,T}} (\phi_a(s)\underline{h})g(s)\mathrm{d}s \\ &+ (1-\pi)\int_{S_{2,T}}^{\infty} (\phi_a(s)\bar{h})g(s)\mathrm{d}s + \pi \int_{S_{1,T}}^{\infty} (\phi_a(s)\bar{h})g(s)\mathrm{d}s\Bigg] \end{aligned}$$

For Case 1, this yields the following optimality conditions that implicitly characterize the cut-off skill types:

$$[S_{1,T}] : \frac{(1-\pi)\cdot(F_{N_{a,T+1}}\cdot\phi(S_{1,T}) - F_{N_{r,T+1}})\cdot(\bar{h}/\underline{h})}{F_{N_{r,T}} - F_{N_{a,T}}\cdot\phi(S_{1,T})} = R$$

$$[S_{2,T}] : \frac{(1-\pi)\cdot(F_{N_{a,T+1}}\cdot\phi(\tilde{S}_{2,T}) - F_{N_{r,T+1}})\cdot(\bar{h}/\underline{h})}{F_{N_{r,T}}\cdot(\bar{h}/\underline{h}) - F_{N_{a,T}}\cdot\phi(\tilde{S}_{2,T})} = R, \quad S_{2,T} = \min\{\tilde{S}_{2,T}, \mathbf{S}\}$$

$$[S_{1,T+1}] : F_{N_{r,T+1}} = F_{N_{a,T+1}}\cdot\phi(S_{1,T+1})$$

$$[S_{2,T+1}] : F_{N_{r,T+1}}\cdot(\bar{h}/\underline{h}) = F_{N_{a,T+1}}\cdot\phi(\tilde{S}_{2,T+1}), \quad S_{2,T+1} = \min\{\tilde{S}_{2,T+1}, \mathbf{S}\}$$

Focus on the first condition, i.e. for $S_{1,T}$. Intuitively, choosing the optimal cut-offs in period T is an investment decision. Putting an additional household to work in the abstract occupation reduces output today, a cost which appears in the denominator (remember that the denominator is zero for $\phi(\mathbf{S})$ without technological change and becomes positive for $S_{1,T} < \mathbf{S}$, see equation (7)). This cost must be worth the benefit of having an additional experienced abstract worker in period $T+1$, which appears in the numerator.

The first condition has a solution with $S_{1,t} < \mathbf{S}$ whenever $\omega \equiv \left(\frac{F_{N_{a,T+1}}}{F_{N_{r,T+1}}}\right) / \left(\frac{F_{N_{a,T}}}{F_{N_{r,T}}}\right) > 1$, i.e. the marginal product of labor in the abstract occupation grows faster in the abstract than in the routine occupation.[30] I assume in the following that this is satisfied. Note that the assumption $\omega > 1$ is a bit stronger than the second bullet of Assumption 1, i.e. that the abstract marginal product grows faster in q than the routine marginal product, holding fixed factor inputs. Intuitively, $\omega > 1$ requires that the relative increase in the marginal product in the abstract occupation driven by $q \uparrow$ is not offset by workers moving from the routine to the abstract occupation. Note that this assumption is satisfied both empirically and in the competitive equilibrium of the quantitative model in the main text, where abstract (and manual) wages grow faster than routine wages despite occupational switching away from the routine occupation. Moreover, I show analytically in Appendix A.2.5 that this assumption is satisfied in the case of a CES production function and no human capital, except for the extreme case of perfect complementarity of labor inputs.
For completeness, in Case 2, $S_{2,T}$ is instead characterized by

$$[S_{2,T}] : \ \frac{(1-\pi) \cdot F_{N_{a,T+1}} \cdot \phi(\tilde{S}_{2,T}) \cdot (\bar{h}/\underline{h} - 1)}{F_{N_{r,T}} \cdot (\bar{h}/\underline{h}) - F_{N_{a,T}} \cdot \phi(\tilde{S}_{2,T})} = R, \quad S_{2,T} = \min\{\tilde{S}_{2,T}, \mathbf{S}\}$$

while the remaining cut-offs are as in Case 1. In Case 3, $S_{1,T}$ is instead characterized by

$$[S_{1,T}] : \ \frac{(1-\pi) \cdot F_{N_{a,T+1}} \cdot \phi(S_{1,T}) \cdot (\bar{h}/\underline{h} - 1)}{F_{N_{r,T}} - F_{N_{a,T}} \cdot \phi(S_{1,T})} = R$$

while the remaining cut-offs are as in Case 2. For future reference, define C_t^*, Y_t^* and ω^* as the values that characterize the optimal allocations arising from the planner's solution in periods $t = T, T+1$.

[30] To see this, rewrite to obtain $\frac{R}{(1-\pi)\bar{h}/\underline{h}} = -\frac{F_{N_{r,T+1}}}{F_{N_{r,T}}} \frac{\omega\phi(S_{1,T})F_{N_{a,T}}/F_{N_{r,T}} - 1}{\phi(S_{1,T})F_{N_{a,T}}/F_{N_{r,T}} - 1}$.

A.2.3 Competitive equilibrium

Next I show that the same cut-off skill levels that I have just derived as the optimal solution to the planner's problem arise from a decentralized equilibrium in which households can freely borrow against future income and in which a representative firm maximizes profits. Note that I allow here already for household heterogeneity by assuming that households differ by skill type. I will show that this heterogeneity by itself does not lead to inefficient occupational choices.

Suppose first that $\forall t, q_t = q$ such that both wages are constant over time. In this case, it is optimal for households to choose the occupation in which they are most productive and work there their entire life. Hence, absent wage growth the cut-off $\mathbf{S}$ from the planner's problem again separates the skill space into routine and abstract workers and the allocation of labor is efficient.

Let me therefore turn to the more interesting case of solving the problem with technological change, $q_{T+1} > q_T$. A household of type (s, j_{T-1}) in period T solves:[31]

$$\max_{\{j_t, c_t\}_{t=T,T+1}} \quad u(c_T) + \hat{\beta} \cdot u(c_{T+1})$$

subject to:

$$c_T + \frac{c_{T+1}}{R/(1-\pi)} = y_T(s, j_{T-1}, j_T) + \frac{y_{T+1}(s, j_T, j_{T+1})}{R/(1-\pi)} \equiv y_L(s, j_{T-1}, j_T, j_{T+1})$$

In a competitive equilibrium, furthermore, a representative firm maximizes profits, implying $w_{j,t} = F_{N_{j,t}}$ and the goods market clears, i.e. aggregate consumption equals output of the representative firm.

Focus now on the occupational choice of an inexperienced household in period $T-1$, $j_{T-1} = \emptyset$. Denoting by $w_{j,t}$ the wage a household earns when working in occupation j, her lifetime income y_L is

$$y_L(s, j_{T-1} = \emptyset, j_T, j_{T+1}) = \begin{cases} \underline{h} \cdot \phi_r(s) \cdot w_{r,T} + \frac{\bar{h} \cdot \phi_r(s) \cdot w_{r,T+1}}{R/(1-\pi)} & \text{if } j_T = r, j_{T+1} = r \\ \underline{h} \cdot \phi_r(s) \cdot w_{r,T} + \frac{\underline{h} \cdot \phi_a(s) \cdot w_{a,T+1}}{R/(1-\pi)} & \text{if } j_T = r, j_{T+1} = a \\ \underline{h} \cdot \phi_a(s) \cdot w_{a,T} + \frac{\bar{h} \cdot \phi_a(s) \cdot w_{a,T+1}}{R(1-\pi)} & \text{if } j_T = a, j_{T+1} = a \end{cases} \tag{8}$$

I want to determine the skill type $S^*_{1,T}$ that is exactly indifferent between working in the routine and the abstract occupation in period T.

[31] I assume that households may die with positive or negative asset holdings. The foreign creditor realizes this and augments the gross interest rate households have to pay by a factor $\frac{1}{1-\pi}$. This way, aggregate debt can be honored in the last period since $R \int s_T = (1-\pi) \int (c_{T+1} - w_{T+1})$, where $s_T = w_T - c_T$.

Case A: (corresponds to Cases 1 and 2 of the planner)

$$\bar{h} \cdot \phi_r(S^*_{1,T}) \cdot w_{r,T+1} > \underline{h} \cdot \phi_a(S^*_{1,T}) \cdot w_{a,T+1}$$

In this case, the skill type $S^*_{1,T}$ is determined by:

$$\underline{h} \cdot \phi_r(S^*_{1,T}) \cdot w_{r,T} + \frac{\bar{h} \cdot \phi_r(S^*_{1,T}) \cdot w_{r,T+1}}{R/(1-\pi)} = \underline{h} \cdot \phi_a(S^*_{1,T}) \cdot w_{a,T} + \frac{\bar{h} \cdot \phi_a(S^*_{1,T}) \cdot w_{a,T+1}}{R/(1-\pi)}$$
$$\Rightarrow \quad \frac{(1-\pi) \cdot (w_{a,T+1} \cdot \phi(S^*_{1,T}) - w_{r,T+1}) \cdot (\bar{h}/\underline{h})}{w_{r,T} - \phi(S^*_{1,T}) \cdot w_{a,T}} = R$$

Case B: (corresponds to Case 3 of the planner)

$$\bar{h} \cdot \phi_r(S^*_{1,T}) \cdot w_{r,T+1} \leq \underline{h} \cdot \phi_a(S^*_{1,T}) \cdot w_{a,T+1}$$

In this case, the skill type $S^*_{1,T}$ is determined by the following equation:

$$\underline{h} \cdot \phi_r(S^*_{1,T}) \cdot w_{r,T} + \frac{\underline{h} \cdot \phi_a(S^*_{1,T}) \cdot w_{a,T+1}}{R/(1-\pi)} = \underline{h} \cdot \phi_a(S^*_{1,T}) \cdot w_{a,T} + \frac{\bar{h} \cdot \phi_a(S^*_{1,T}) \cdot w_{a,T+1}}{R/(1-\pi)}$$
$$\Rightarrow \quad \frac{(1-\pi) \cdot w_{a,T+1} \cdot \phi(S^*_{1,T}) \cdot (\bar{h}/\underline{h} - 1)}{w_{r,T} - \phi(S^*_{1,T}) \cdot w_{a,T}} = R$$

Since the optimal behavior of the representative firm ensures that wages equal the marginal product of labor, $w_{j,t} = F_{N_{j,t}}$, the condition that characterizes $S^*_{1,T}$ coincides with that for $S_{1,T}$ derived in the planner's problem, so the two cut-offs are equal. The same property can be shown in an analogous fashion for the cut-off $S_{1,T+1}$ (using the problem of inexperienced workers in $T+1$) as well as for the cut-offs $S_{2,T}$ and $S_{2,T+1}$ (using the problem of experienced routine workers with $j_{T-1} = r$ and $j_T = r$, respectively).

I have proven, therefore, that the allocation of labor in the decentralized economy is efficient and coincides with the first-best allocation chosen by the social planner. Given the heterogeneity in skill types (and strictly concave utility), however, consumption levels inefficiently differ across households, being lower (higher) for less (more) skilled households. This can easily be remedied, however, by introducing an appropriate (type-specific) lump-sum tax scheme which leaves occupational choices unaffected but ensures $c(s, j_{t-1}) = C^*_t$ for all types (s, j_{t-1}) and periods $t = T, T+1$.[32] Hence, in this slightly adjusted version of the model with lump-sum transfers the

[32]The necessary transfer is conditional on skill type and on the original occupation in $T-1$ (so as not to distort individual occupational choices going forward). In T it is: $tr(s, j_{T-1}) = Y^*_T - y_T(s, j_{T-1}, j^*_T)$ and in $T+1$ it is $tr(s, j^*_T(j_{T-1})) = Y^*_{T+1} - y_{T+1}(s, j^*_T(j_{T-1}), j^*_{T+1}(j^*_T(j_{T-1})))$, where j^*_t denotes the optimal occupational choice in period t, explicitly derived for the case of newborn workers above. This transfer cancels out in aggregate and ensures $c(s, j_{t-1}) = C^*_t$ for all types (s, j_{t-1}) and periods $t = T, T+1$.

welfare theorems hold: the market equilibrium is Pareto optimal and the planner's solution can be decentralized. In contrast, the introduction of a borrowing constraint in the next subsection will lead to a departure from the first-best allocation. As in the simpler model in Section 2, the borrowing constraint inefficiently distorts individual occupational choices.

For future reference, let me define wage growth in occupation j as

$$\omega_j \equiv \frac{w_{j,T+1}}{w_{j,T}} \text{ and, similar to above, } \omega \equiv \frac{\omega_a}{\omega_r}$$

and let me explicitly solve for $\phi(S_{1,T})$:

Case A:

$$\phi(S_{1,T}) = \frac{w_{r,T}}{w_{a,T}} \cdot \frac{1 + \frac{1-\pi}{R} \cdot (\bar{h}/\underline{h}) \cdot \omega_r}{1 + \frac{1-\pi}{R} \cdot (\bar{h}/\underline{h}) \cdot \omega_a}$$

Case B:

$$\phi(S_{1,T}) = \frac{w_{r,T}}{w_{a,T}} \cdot \frac{1}{1 + \frac{1-\pi}{R} \cdot \omega_a \cdot (\bar{h}/\underline{h} - 1)}$$

A.2.4 Competitive equilibrium with hand-to-mouth households

I now analyze the equilibrium allocation when all households in the economy are hand-to-mouth, i.e. they consume their current income in both periods. Consuming current income would endogenously arise as the optimal choice of all households if wages grew over time in both occupations and households held zero assets and were exogenously prevented from borrowing against future income.[33]

Suppose first that $\forall t, q_t = q$ such that both wages are constant over time. In this case, it is optimal for hand-to-mouth households to choose the occupation in which they are most productive and work there their entire life. Hence, absent wage growth the cut-off $\mathbf{S}$ from the planner's problem again separates the skill space into routine and abstract workers and the allocation of labor is efficient, as in the case without borrowing constraints. I therefore turn to analyzing the problem with technological change, i.e. $q_{T+1} > q_T$. The problem of a household with type (s, j_{T-1}) is:

$$\max_{j_T, j_{T+1}} u(y_T) + \hat{\beta} \cdot u(y_{T+1})$$

Note that occupational choices lead to the same implications for income y_t as shown in (8).

To make progress, I assume log utility in what follows, i.e. $u(c) = \ln(c)$. As before, focus on the decision of an inexperienced household, i.e. with $j_{T-1} = \emptyset$.

[33]Under these assumptions, the proof of this is analogous to the one provided in Appendix A.1.2 for the simpler model.

Case A:

$$\bar{h} \cdot \phi_r(S^+_{1,T}) \cdot w_{r,T+1} > \underline{h} \cdot \phi_a(S^+_{1,T}) \cdot w_{a,T+1}$$

In this case, $S^+_{1,T}$ is determined by:

$$(\phi_r(S^+_{1,T}) \cdot w_{r,T} \cdot \underline{h}) \cdot (\phi_r(S^+_{1,T}) \cdot w_{r,T+1} \cdot \bar{h})^{\hat{\beta}} = (\phi_a(S^+_{1,T}) \cdot w_{a,T} \cdot \underline{h}) \cdot (\phi_a(S^+_{1,T}) \cdot w_{a,T+1} \cdot \bar{h})^{\hat{\beta}}$$

$$\phi(S^+_{1,T}) = \frac{w_{r,T}}{w_{a,T}} \cdot \left(\frac{\omega_r}{\omega_a}\right)^{\frac{\hat{\beta}}{1+\hat{\beta}}}$$

Case B:

$$\bar{h} \cdot \phi_r(S^+_{1,T}) \cdot w_{r,T+1} \leq \underline{h} \cdot \phi_a(S^+_{1,T}) \cdot w_{a,T+1}$$

In this case, the skill type $S^+_{1,T}$ that is indifferent between working in the routine and the abstract occupation in period T is determined by:

$$\begin{aligned} \ln\Big(\phi_r(S^+_{1,T}) \cdot w_{r,T} \cdot \underline{h}\Big) + \hat{\beta} \cdot \ln\Big(\phi_a(S^+_{1,T}) \cdot w_{a,T+1} \cdot \underline{h}\Big) \\ = \ln\Big(\phi_a(S^+_{1,T}) \cdot w_{a,T} \cdot \underline{h}\Big) \\ + \hat{\beta} \cdot \ln\Big(\phi_a(S^+_{1,T}) \cdot w_{a,T+1} \cdot \bar{h}\Big) \end{aligned}$$

$$(\phi_r(S^+_{1,T}) \cdot w_{r,T} \cdot \underline{h}) \cdot (\phi_a(S^+_{1,T}) \cdot w_{a,T+1} \cdot \underline{h})^{\hat{\beta}} = (\phi_a(S^+_{1,T}) \cdot w_{a,T} \cdot \underline{h}) \cdot (\phi_a(S^+_{1,T}) \cdot w_{a,T+1} \cdot \bar{h})^{\hat{\beta}}$$

$$\phi(S^+_{1,T}) = \frac{w_{r,T}}{w_{a,T}} \cdot (\bar{h}/\underline{h})^{-\hat{\beta}}$$

To show that fewer workers choose to work in the abstract occupation in T when borrowing constraints bind than when borrowing is unrestricted I next prove that $\phi(S^+_{1,T}) > \phi(S^*_{1,T})$, which by monotonicity of $\phi(s)$ implies $S^+_{1,T} > S^*_{1,T}$.

Case A:

$$\phi(S^*_{1,T}) < \phi(S^+_{1,T})$$

$$\frac{1 + \frac{1-\pi}{R} \cdot (\bar{h}/\underline{h}) \cdot \omega_r}{1 + \frac{1-\pi}{R} \cdot (\bar{h}/\underline{h}) \cdot \omega_a} < \left(\frac{\omega_r}{\omega_a}\right)^{\frac{\hat{\beta}}{1+\hat{\beta}}}$$

Remembering $\omega \equiv \frac{\omega_a}{\omega_r}$ and taking the log on both sides,

$$0 < \underbrace{\ln\Big(1 + \frac{1-\pi}{R} \cdot (\bar{h}/\underline{h}) \cdot \omega_r \cdot \omega\Big) - \ln\Big(1 + \frac{1-\pi}{R} \cdot (\bar{h}/\underline{h}) \cdot \omega_r\Big) - \frac{\hat{\beta}}{1+\hat{\beta}} \ln(\omega)}_{H(\omega)}$$

For this inequality to be fulfilled it needs to hold that $\forall \omega > 1, H(\omega) > 0$. Note that $H(1) = 0$. Hence it suffices to show that $\forall \omega \geq 1, H'(\omega) > 0$.

$$
\begin{aligned}
H'(\omega) = \frac{\frac{1-\pi}{R} \cdot (\bar{h}/\underline{h}) \cdot \omega_r}{1 + \frac{1-\pi}{R} \cdot (\bar{h}/\underline{h}) \cdot \omega_r \cdot \omega} - \frac{\hat{\beta}}{1+\hat{\beta}} \frac{1}{\omega} &> 0 \\
\frac{\frac{1-\pi}{R} \cdot (\bar{h}/\underline{h}) \cdot \omega_r}{\frac{1}{\omega} + \frac{1-\pi}{R} \cdot (\bar{h}/\underline{h}) \cdot \omega_r} &> \frac{\hat{\beta}}{1+\hat{\beta}} \\
\frac{1+\hat{\beta}}{\hat{\beta}} &> 1 + \frac{1}{\omega \cdot (\bar{h}/\underline{h}) \cdot \frac{1-\pi}{R} \cdot \omega_r} \\
1 &> \frac{\hat{\beta}}{\omega \cdot (\bar{h}/\underline{h}) \cdot \frac{1-\pi}{R} \cdot \omega_r} \\
\underbrace{\omega_a \cdot (\bar{h}/\underline{h})}_{>1} &> \underbrace{\beta \cdot R}_{\leq 1}
\end{aligned}
$$

Case B:

$$
\begin{aligned}
\phi(S^*_{1,T}) &< \phi(S^+_{1,T}) \\
\frac{1}{1 + \frac{1-\pi}{R} \cdot \omega_a \cdot (\bar{h}/\underline{h} - 1)} &< (\bar{h}/\underline{h})^{-\hat{\beta}} \\
0 &< \underbrace{1 + \frac{1-\pi}{R} \cdot \omega_a \cdot (\bar{h}/\underline{h} - 1) - (\bar{h}/\underline{h})^{\hat{\beta}}}_{\Lambda(\bar{h}/\underline{h})}
\end{aligned}
$$

For this inequality to be fulfilled it needs to hold that $\forall (\bar{h}/\underline{h}) > 1, \Lambda(\bar{h}/\underline{h}) > 0$. Note that $\Lambda(1) = 0$. Hence it suffices to show that $\forall (\bar{h}/\underline{h}) \geq 1, \Lambda'(\bar{h}/\underline{h}) > 0$.

$$
\begin{aligned}
\Lambda'(\bar{h}/\underline{h}) = \frac{1-\pi}{R} \cdot \omega_a - \hat{\beta}(\bar{h}/\underline{h})^{\hat{\beta}-1} &> 0 \\
\underbrace{\omega_a}_{\geq 1} &> \underbrace{\beta \cdot R}_{\leq 1} \cdot \underbrace{(\bar{h}/\underline{h})^{\hat{\beta}-1}}_{<1}
\end{aligned}
$$

This completes the proof that $S^+_{1,T} > S^*_{1,T} = S_{1,T}$ and the same result can be shown to hold for the cut-off $S_{2,T}$ when analyzing the problem of an experienced routine worker. Hence, fewer skill types work in the abstract occupation in T when households are hand-to-mouth than is optimal. This in turn leads to a shortage of experienced abstract workers in period $T + 1$.

A.2.5 CES production and no human capital

With CES production, we have

$$\begin{aligned} F(N_a, N_r; q) &= \left(qN_a^{\frac{\sigma-1}{\sigma}} + N_r^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}} \\ \Rightarrow F_{N_a} &= qF^{\frac{1}{\sigma}} N_a^{-\frac{1}{\sigma}} \\ \Rightarrow F_{N_r} &= F^{\frac{1}{\sigma}} N_r^{-\frac{1}{\sigma}} \end{aligned}$$

Therefore,

$$\frac{F_{N_a}}{F_{N_r}} = q\left(\frac{N_a}{N_r}\right)^{-\frac{1}{\sigma}} \tag{9}$$

With no human capital,[34] labor as a function of the cut-off S is

$$N_a = \int_S^\infty \phi_a(s)g(s)\mathrm{d}s, \quad N_r = \int_{-\infty}^S \phi_r(s)g(s)\mathrm{d}s$$

Output maximization leads to a condition that implicitly characterizes the cut-off $\mathbf{S}$ in the efficient allocation:

$$\frac{1}{\phi(\mathbf{S})} = q\left(\frac{N_a(\mathbf{S})}{N_r(\mathbf{S})}\right)^{-\frac{1}{\sigma}} \tag{10}$$

When q rises, the only way for this equation to hold is for $\mathbf{S}$ to fall (reallocation from routine to abstract labor), thereby raising $\frac{1}{\phi(\mathbf{S})}$ and lowering $\left(\frac{N_a(\mathbf{S})}{N_r(\mathbf{S})}\right)^{-\frac{1}{\sigma}}$. Note, however, that since $\phi(\mathbf{S})$ falls (it is monotonically increasing in $\mathbf{S}$), the rise in q is larger than the fall in $\left(\frac{N_a(\mathbf{S})}{N_r(\mathbf{S})}\right)^{-\frac{1}{\sigma}}$, for all $\sigma > 0$. This proves that the marginal product of abstract relative to routine labor (9) increases in q, even after taking into account the efficient reallocation of labor.

B Data Sources and Measurement

B.1 Labor market polarization and broad occupational groups

To describe the phenomenon of labor market polarization, Autor and Dorn (2013) order 318 detailed occupations according to their skill level, which they proxy by the average hourly wage earned in the occupation in 1980. They show that between 1980 and 2005 occupations at the bottom and

[34] Equivalently, one can think of this example as a case with human capital but certain death after one period, i.e. $\pi = 1$.

at the top of the skill distribution gained both in terms of employment shares and in wages relative to occupations in the center of the distribution. Scrutinizing the task content of each occupation using the US Department of Labor's Dictionary of Occupational Titles, they divide occupations into six broad groups. For tractability, the literature often subsumes these further into three broad occupational groups: manual, routine and abstract (Jaimovich and Siu, 2020; vom Lehn, 2020).[35]

The first group of occupations are intensive in **manual** tasks, i.e. they require eye-hand-foot coordination, adapting to new surroundings often, and (non-trivial) interaction with other humans. These occupations can be found at the bottom of the wage distribution and are typically service occupations. Examples for manual occupations are health and nursing aides, workers in child and elderly care, bus drivers, waiters and waitresses, door-to-door/street sales, or janitors and gardeners.

The second group consists of occupations that require **routine** tasks to be performed, i.e. calculations, record-keeping, repetitive customer service, repetitive assembly, picking or sorting. Many occupations in the middle of the wage distribution belong to this group. Typical examples of routine occupations are bookkeepers, accounting clerks, secretaries, bank tellers, as well as machine operators and assemblers or butchers and meat cutters.

Occupations of the third group require **abstract** tasks, i.e. cognitive thinking, forming/testing hypotheses, persuading, managing or organizing. These kind of jobs are found in the high-wage occupations, and typical examples are school teachers, managers, police and detectives, public service (e.g. city planners), or engineers.

The term "polarization" refers to the fact that the occupational wage and employment distributions have shifted towards their poles since the 1980s, while the middle of these distributions, where many of the routine occupations are found, has "hollowed out". The employment share in the routine occupations fell from 58.5% in 1980 to 46.2% in 2005, while it increased from 31.6% to 40.9% in the abstract, and from 9.9% to 12.9% in the manual occupations (Autor and Dorn, 2013). One commonly cited cause for the decline of the routine occupations is technological change that substitutes for routine labor (Autor, Levy, et al., 2003; Jaimovich, Saporta-Eksten, et al., 2021; vom Lehn, 2020). The fact that capital, most prominently ICT capital, has become much cheaper over the recent decades has led to a replacement of tasks formerly performed by humans with machines. Manual and abstract jobs are not so easily substituted by

[35]Following Jaimovich and Siu (2020), the services occupations are the only occupations I categorize as manual. See Appendix B for details. The assignment is complete and mutually exclusive: each of the detailed occupations is considered to be part of exactly one of the three broad groups.

ICT capital, as they require either non-trivial interaction with humans or tasks such as managing and organizing, all of which machines struggle to excel at. Following the literature, a falling relative price of ICT capital drives technological change in the model of Section 4.

B.2 Classification of broad occupational groups

Based on earlier work by Dorn (2009), Autor and Dorn (2013) map Census Occupation Codes from 1950 to 2005 into a time-consistent set of occupations. They then form six broad occupational groups, based on the task content in each occupation which they in turn derive from the US Department of Labor's Document of Occupational Titles. Throughout this paper, I use the following categorization of these six groups into manual, routine and abstract occupations, which coincides with the concepts used in Jaimovich and Siu (2020) and in the appendix of vom Lehn (2020):

- Abstract:
  - management/professional/technical/financial sales/public security occupations
- Routine:
  - administrative support and retail sales occupations
  - precision production and craft occupations
  - machine operators, assemblers and inspectors
  - transportation/construction/mechanics/mining/agricultural occupations
- Manual:
  - low-skill services

B.3 PSID

B.3.1 Sample selection

I use waves 1976 to 2017 of the PSID and I follow Cortes (2016) in terms of sample selection and when defining variables. In particular, Cortes (2016) uses only the core PSID sample, drops all individuals who have never been a household head or a wife and only keeps observations aged 16–64. Note that while Cortes (2016) uses data only up until 2007, I include data until 2017, i.e. five more waves of the PSID. Whenever using real wages (or changes thereof) in the analysis, I follow Cortes (2016) in excluding observations with log hourly real wages below 1.10$ or above 54.60$ in terms of 1979 dollars (3.87$ and 192.30$ in terms of 2019 dollars). This is relevant for the results in Section 3.1 and when estimating returns to tenure in the calibration. Like Cortes (2016), I use the Consumer Price Index for All Urban Consumers: All Items in U.S. City Average (CPIAUCSL, downloaded from the Federal Reserve Economic Database) to deflate nominal quantities. The parts of the empirical analysis that are based on PSID data (Sections 3.1 and 3.3) further only uses male observations to make the results comparable to Cortes (2016)'s.

As mentioned in the main text, in order to be consistent throughout this paper I marginally deviate from Cortes (2016) when assigning the broad occupation to individuals. While Cortes (2016) groups occupations according to the classification given in Acemoglu and Autor (2011), I form groups from the six broad groups listed in Autor and Dorn (2013) as described above. The two concepts yield very similar results, and Cortes (2016) shows in his appendix that using a classification that is closer to the one used by Autor and Dorn (2013) leads to very similar results as his baseline.

B.3.2 Wage changes of switchers vs stayers

To remind the reader: To obtain estimates of how wages of switchers from routine to abstract and manual occupations evolved relative to routine stayers at horizon h, Cortes (2016) performs the following regression

$$\Delta_h \ln(wage_{it}) = \beta_m^h \cdot D_{imt} + \beta_a^h \cdot D_{iat} + \gamma_t^h + u_{iht}$$

where $\Delta_h \ln(wage_{it})$ is the change in the log real hourly wage of individual i between years t and $t+h$, D_{ijt} is a dummy variable equal to zero if the individual was working in the routine occupation in both t and $t+1$ (or t and $t+2$ for $h \geq 2$) and one if she switched from the routine occupation to occupation j, γ_t captures time fixed effects, and u_{iht} is a mean zero error term. The solid blue lines in Figure 7 visualize his estimates of β_m^h (left panel) and β_a^h (right panel), which I already printed in Table 1 in the main text. The dotted blue lines show my own estimates, which employ the slightly different occupational groups as well as the additional five waves of the PSID. As can be seen, the estimates are nearly identical to Cortes (2016)'s.

The grey dotted lines show the estimates when deleting all observations from the sample who see immediate wage gains (i.e. between t to $t+1$) when leaving the routine occupation. Since the PSID went to bi-annual frequency after 1997, so that I do not observe year-on-year wage changes in the later waves, I restrict the sample to observations prior to 1997 for

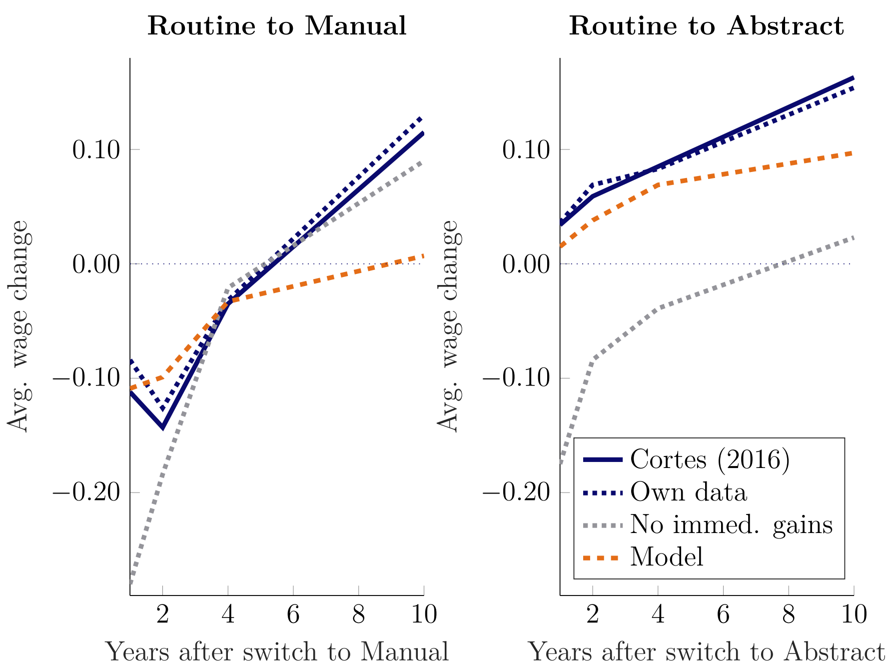

Figure 7: Average log hourly real wages of switchers from routine to manual (left) and to abstract (right) compared to routine stayers.
Notes: Blue solid lines are the estimates from Cortes (2016, Table 3). "Own data" replicates his analysis. "No immed. gains" excludes all individuals from the sample who saw positive wage changes from leaving the routine occupation on impact. For obtaining the model values I perform the same regression as Cortes (2016), only using data from a synthetic panel of 10,000 households, simulated between 1980 and 2020 from the full model of Section 4.

computing these estimates. Estimating the wage effects only on workers who see initial wage losses when leaving the routine occupations obviously lowers the average wage gains obtained from switching, but, as can be seen, even this subset of workers saw faster wage growth than the stayers. The dashed red lines show the estimates from the synthetic panel simulated from the full model of Section 4. While the initial average wage change (i.e. after year 1) is targeted in the calibration, the remaining horizons are untargeted, and still provide a relatively good fit to Cortes (2016)'s estimates. Long-run wage growth of switchers is, however, somewhat lower in the model than in the data, especially for switches to the manual occupation (see discussion in Appendix D).

B.3.3 Analogues of Table 1 for abstract and manual workers

Tables 4 and 5 are the analogues of Table 1 for abstract and manual workers, respectively.

B.3.4 Definition of being hand-to-mouth

In classifying households as either being hand-to-mouth or not I follow Kaplan, Violante, et al. (2014). For a detailed record of how they define income, liquid and illiquid wealth, see pages 122-3 in Kaplan, Violante, et

Table 4: Wage changes for abstract workers relative to stayers, according to direction of switch

Change in real wages between year t and year ...	To routine occupation	To manual occupation
$t+1$	−4.7% (−4.0%)	−6.8% (−11.3%)
$t+2$	−7.9% (−7.9%)	−9.0% (−11.2%)
$t+4$	−9.8% (−6.4%)	−3.8% (−6.0%)
$t+10$	−20.3% (−5.1%)	−18.5% (−10.7%)

Notes: Workers who stay in abstract occupations are the omitted category. Numbers show change in real wages between year t and year $t+x$, compared to workers who stayed in abstract occupations reported in Cortes (2016, Table 5). Numbers in parentheses are estimates using my own sample.

Table 5: Wage changes for manual workers relative to stayers, according to direction of switch

Change in real wages between year t and year ...	To routine occupation	To abstract occupation
$t+1$	10.6% (10.9%)	9.5% (12.6%)
$t+2$	6.5% (6.8%)	15.4% (15.4%)
$t+4$	12.5% (9.6%)	23.5% (16.1%)
$t+10$	1.8% (12.4%)	48.9% (28.7%)

Notes: Workers who stay in manual occupations are the omitted category. Numbers show change in real wages between year t and year $t+x$, compared to workers who stayed in manual occupations reported in Cortes (2016, Table 5). Numbers in parentheses are estimates using my own sample.

al. (2014). I use information on income and wealth on determining whether a household is hand-to-mouth as they do and I detail this procedure in Appendix B.5.

B.3.5 Returns to tenure

As mentioned in the main text, when estimating returns to tenure in the calibration, I first follow the procedure outlined in Cortes (2016, Section VI C). In particular, using PSID data from 1981 to 2017, I estimate the following equation:

$$\mathcal{Y}_{it} = \sum_j D_{ijt} \left(\beta_{j1} Ten_{ijt} + \beta_{j2} Ten_{ijt}^2 + \gamma_{ij} \right) + \delta X_{it} + u_{it}$$

where $\mathcal{Y}_{it}$ is the log real wage of individual i at time t, D_{ijt} is a dummy indicating whether the individual was working in occupation $j \in \{m, r, a\}$ at time t, Ten_{ijt} is occupational tenure, γ_{ij} is an occupation fixed effect for each individual, and X_{it} are controls. I only deviate from Cortes (2016) by leaving females in the sample (as employment shares used in the calibration also relate to both males and females (Autor and Dorn, 2013)), and by only assigning a value for broad occupational tenure to individuals once I have

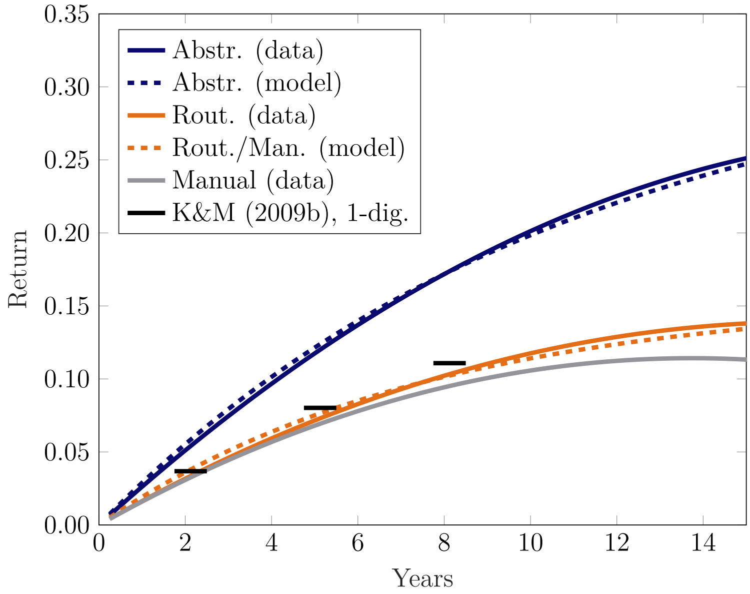

Figure 8: Returns to occupational tenure, estimated from PSID data (1981-2017).

Notes: I estimate $\mathcal{Y}_{it} = \sum_j D_{ijt} \left(\beta_{j1} Ten_{ijt} + \beta_{j2} Ten_{ijt}^2 + \gamma_{ij}\right) + \delta X_{it} + u_{it}$, where $\mathcal{Y}_{it}$ is the log real wage of individual i at time t, D_{ijt} is a dummy indicating whether the individual was working in occupation $j \in \{m, r, a\}$ at time t, Ten_{ijt} is occupational tenure, γ_{ij} is an occupation fixed effect for each individual, and X_{it} are controls (unionization and marital status, region of residence, year dummies and year-occupation dummies). I then calibrate the parameters of the Poisson process in the model $(\bar{h}_j - \underline{h})$ and λ_j^h to obtain the closest fit (in a mean squared sense) to the data over the first 20 years.

observed them making an occupational switch. Following Kambourov and Manovskii (2009b), every time a person indicates working in a different occupational group than in her most previous report, Cortes (2016) considers the person to have made a switch and occupational tenure is reset. For individuals who are observed for the first time in the survey, Cortes (2016) constructs estimates of occupational tenure by setting it equal to the person's stated tenure at her employer or in her current position. This is of course an imperfect estimate of tenure in the broad occupation and adds significant measurement error to the variable. As I have many additional observations in comparison to Cortes (2016) because of the five more waves of the PSID that I use, I choose to limit myself to individuals whom I have observed to make a switch and for whom I can therefore construct the measure relatively cleanly.

The results are shown in Figure 8. In line with previous literature (Sullivan, 2010), I find that returns to tenure are higher in the abstract than in the other two occupations. Moreover, I do not find statistically significant differences between β_{m1} and β_{r1}, nor between β_{m2} and β_{r2}. I therefore use the estimates of β_{j1} and β_{j2} and approximate the tenure profiles once for the abstract and once jointly for the routine and for the manual occupation.

Table 6: Switching decision, estimated with interaction terms.

	(1)	(2)	(3)	(4)
HtM=0 × wage_gain=1	0.025	0.026*	0.098	0.11*
	(0.016)	(0.016)	(0.061)	(0.062)
HtM=1 × wage_gain=0	-0.039**	-0.031	-0.16*	-0.13
	(0.020)	(0.020)	(0.087)	(0.089)
HtM=1 × wage_gain=1	-0.0024	0.0083	-0.013	0.029
	(0.017)	(0.017)	(0.066)	(0.067)
Occupational tenure	-0.014***	-0.014***	-0.083***	-0.079***
	(0.00098)	(0.00095)	(0.0072)	(0.0071)
Skill		0.090***		0.33***
		(0.012)		(0.045)
Controls	Yes	Yes	Yes	Yes
Observations	4882	4882	4882	4882
Model	OLS	OLS	Probit	Probit

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Data are from the PSID, years 1999–2017. Sample selection is as in Cortes (2016). Definition of HtM status is as in Kaplan, Violante, et al. (2014). Dependent variable is whether or not individual leaves routine for the manual occ. between t and $t+2$. Skill is a dummy for whether the individual has received more than 12 years of education. Occ. tenure is years of uninterrupted tenure in broad occupation. Additional controls are: region dummies, age, age squared, unionization status, married status, year. Standard errors are clustered at the individual level. Columns with probit model show probit coefficient.

B.4 Extensions for Section 3.3

B.4.1 Interaction terms

In Table 6 I split up the independent variable "HtM" into an interaction term between being hand-to-mouth and experiencing wage gains upon switching. Consistent with the intuition from the quantitative model it is especially those switches that entail wage losses (second row) that hand-to-mouth agents are especially likely to avoid. Note that since average marginal effects are not well-defined for levels of interaction terms, I display the coefficients of the probit model in this case.

B.4.2 Liquid assets and switching out of abstract and manual occupation

Next, I ask whether being hand-to-mouth also predicts switches out of abstract (Table 7) and out of manual (Table 8) occupations. None of the estimated coefficients on the hand-to-mouth dummy are statistically different from zero. This is in line with the quantitative model of Section 4, in which liquid assets are an especially strong determinant of switching

Table 7: Switching decision, exit from abstract occupations.

	(1)	(2)	(3)	(4)
HtM	0.0025	-0.0013	0.00095	-0.0028
	(0.012)	(0.012)	(0.011)	(0.011)
Occupational tenure	-0.011***	-0.0098***	-0.013***	-0.012***
	(0.00072)	(0.00072)	(0.0011)	(0.0011)
Skill		-0.11***		-0.087***
		(0.015)		(0.012)
Controls	Yes	Yes	Yes	Yes
Observations	5550	5550	5550	5550
Model	OLS	OLS	Probit	Probit

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Data are from the PSID, years 1999–2017. Sample selection is as in Cortes (2016). Definition of HtM status is as in Kaplan, Violante, et al. (2014). Dependent variable is whether or not individual leaves abstract occ. between t and $t+2$. Skill is a dummy for whether the individual has received more than 12 years of education. Occ. tenure is years of uninterrupted tenure in broad occupation. Additional controls are: region dummies, age, age squared, unionization status, married status, year. Standard errors are clustered at the individual level. Columns with probit model show average marginal effect.

behavior only among routine workers.

B.4.3 Preference types

Aguiar et al. (2020) have recently de-emphasized the role of "bad luck", i.e. low productivity draws, in determining an individual's hand-to-mouth status. Rather they argue that there exist certain (fixed) preference types who are especially likely to be hand-to-mouth, while other preference types are less likely to (ever) be hand-to-mouth. I discuss this alternative theory of being hand-to-mouth in the context of the quantitative model and policy analysis in the robustness checks in Appendix F.1.

In the empirical context, Aguiar et al. (2020)'s results could imply that the switching decision is in fact driven by being of a certain preference type (e.g. being very patient/impatient), and that the current hand-to-mouth status is only a proxy for being of a certain preference type. This would imply, however, that a worker's hand-to-mouth status at *any* point in time, regardless of when it was observed, should be an equally good predictor of the switching decision, as it would carry the same information about the worker's preference type as the current status.

To test this, I therefore add the last ever observed hand-to-mouth status of a worker before dropping out of the sample as a control variable in the regression. If being hand-to-mouth is indeed driven by being of a fixed preference type, one would expect the last observed hand-to-mouth status

Table 8: Switching decision, exit from manual occupations.

	(1)	(2)	(3)	(4)
HtM	0.012	0.012	0.0092	0.0088
	(0.030)	(0.030)	(0.029)	(0.029)
Occupational tenure	-0.018***	-0.019***	-0.022***	-0.022***
	(0.0031)	(0.0032)	(0.0043)	(0.0043)
Skill		-0.0037		-0.0036
		(0.029)		(0.029)
Controls	Yes	Yes	Yes	Yes
Observations	1135	1135	1135	1135
Model	OLS	OLS	Probit	Probit

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Data are from the PSID, years 1999–2017. Sample selection is as in Cortes (2016). Definition of HtM status is as in Kaplan, Violante, et al. (2014). Dependent variable is whether or not individual leaves manual occ. between t and $t + 2$. Skill is a dummy for whether the individual has received more than 12 years of education. Occ. tenure is years of uninterrupted tenure in broad occupation. Additional controls are: region dummies, age, age squared, unionization status, married status, year. Standard errors are clustered at the individual level. Columns with probit model show average marginal effect.

to be similarly informative about the switching decision as the current status.[36] As Table 9 shows, however, this is not the case. The coefficient on the current hand-to-mouth status is largely unchanged and the coefficient on the last observed status is not statistically significant. This indicates that there is more information in the current than in some distant hand-to-mouth status. I view this as evidence against the notion that the switching decision in my sample is driven by fixed preference types.

B.4.4 Endogeneity of being hand-to-mouth

Using data from the 2008-09 financial crisis, in this subsection I present suggestive evidence that attempts to identify a causal link between liquid assets and switching decision. In particular, I ask whether the interaction of owning a house and having debt could explain the subsequent switching decision in 2009. The reasoning for this is as follows. As is well-known, house prices dropped significantly during the financial crisis. At the same time, many used their "house as a bank account" in the run-up to the crisis, borrowing heavily against home equity (Mian, Rao, et al., 2013; Mian and Sufi, 2011). Insofar as the drop in house prices was an exogenous event,

[36]The last ever observed hand-to-mouth status is on average more distant in time (four waves) to the observed switching decision of an individual than the first ever observed hand-to-mouth status (two waves). Since I want to isolate the role of fixed preference type and not have the measure for the type be confounded by events close to the observed switch, I opt for using the last observed hand-to-mouth status as a proxy for the preference type.

Table 9: Switching decision and liquid asset holdings (with last ever observed hand-to-mouth status as additional control)

	(1)	(2)	(3)	(4)	(5)	(6)
HtM	-0.029**	-0.023*	-0.031**	-0.026**		
	(0.012)	(0.012)	(0.012)	(0.012)		
Last observed HtM status	-0.0039	0.0079	-0.0018	0.0089	-0.0010	0.0093
	(0.014)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)
Occupational tenure	-0.014***	-0.014***	-0.021***	-0.020***	-0.021***	-0.020***
	(0.00099)	(0.00095)	(0.0018)	(0.0017)	(0.0018)	(0.0017)
Skill		0.090***		0.33***		0.32***
		(0.012)		(0.046)		(0.046)
Poor HtM					-0.046***	-0.034**
					(0.017)	(0.017)
Wealthy HtM					-0.024*	-0.022
					(0.014)	(0.014)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4882	4882	4882	4882	4882	4882
Model	OLS	OLS	Probit	Probit	Probit	Probit

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Data are from the PSID, years 1999–2017. Sample selection is as in Cortes (2016), see also Appendix B.3. Dependent variable is whether or not individual leaves routine occ. between t and $t + 2$. Definition of HtM status is as in Kaplan, Violante, et al. (2014). Last observed HtM status is the last status observed before worker drops out of the sample. Skill is a dummy for whether the individual has received more than 12 years of education. Occ. tenure is years of uninterrupted tenure in broad occupation. Additional controls are: region dummies, age, age squared, unionization status, married status, year. Standard errors are clustered at the individual level. Columns with probit model show average marginal effect.

diminishing homeowners' ability to pay back debt by borrowing against equity, being a homeowner with positive debt could be interpreted as an exogenous shift into being hand-to-mouth in 2009, or a "hand-to-mouth shock".[37]

Table 10 shows that, indeed, when using data from 2009 and regressing the subsequent switching decision on a dummy for being a homeowner, debt, the interaction of the two, as well as the same control variables as in the baseline regressions, the interaction term enters negatively. This is what one would expect given the mechanism I propose: distressed homeowners, no longer being able to borrow against their equity, are up against their borrowing constraint and are less willing to switch the occupation. However, in part owing to the very small sample size (given that I condition on the year being 2009), the results are not always statistically significant (see columns 3 and 4, i.e. the probit regressions). Still, it seems re-assuring that the point estimate is negative, as expected.

B.5 SCF

I use waves 1989 to 2019 of the SCF and I follow Kaplan, Violante, et al. (2014) in terms of sample selection and when defining variables. In particular, I consider all households whose head is aged 22–79, and discard those who report negative labor income, and those whose only positive income stems from self-employment. I refer the reader to Kaplan, Violante, et al. (2014) for further details.

B.5.1 Definition of being hand-to-mouth

I define the liquid asset holdings of a household m_{it} as the sum of cash, money market accounts, checking/savings/call accounts, prepaid cards, directly held stocks, bonds, non-money-market mutual funds, minus revolving consumer debt. Income y_{it} collects labor earnings, regular private transfers (e.g. child support, alimony), and public transfers (e.g. unemployment benefits, food stamps, Social Security Income). Income y_{it} corresponds to bi-weekly income, as this is the most common frequency of payment in the US (Kaplan, Violante, et al., 2014).

Households are considered hand-to-mouth if and only if one of the following two conditions is true:

$$0 \leq m_{it} \leq \frac{y_{it}}{2}$$

[37] Connecting the concepts of being hand-to-mouth and homeownership has been done before. For instance, Cloyne et al. (2020) form three groups in the population, proxying the poor hand-to-mouth by being a renter, the wealthy hand-to-mouth by being a homeowner with a mortgage, and the not-hand-to-mouth by being a homeowner without a mortgage.

Table 10: Switching decision and being a distressed homeowner in 2009

	(1)	(2)	(3)	(4)
Homeowner x Debt	-0.0014**	-0.0016**	-0.0017	-0.0014
	(0.00065)	(0.00068)	(0.0011)	(0.0012)
Homeowner	0.052	0.044	0.046	0.032
	(0.047)	(0.047)	(0.046)	(0.046)
Debt ('000$)	-0.00012	-0.00013	-0.00026	-0.00047
	(0.00012)	(0.00014)	(0.00038)	(0.00065)
Occupational tenure	-0.016***	-0.015***	-0.020***	-0.019***
	(0.0029)	(0.0028)	(0.0058)	(0.0055)
Skill		0.13***		0.48***
		(0.033)		(0.13)
Controls	Yes	Yes	Yes	Yes
Observations	575	575	575	575
Model	OLS	OLS	Probit	Probit

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Data are from the PSID, year 2009. Sample selection is as in Cortes (2016). Dependent variable is whether or not individual leaves routine occ. between t and $t + 2$. Homeowner is a dummy equal to one if the worker has non-zero home equity. Debt is item ER43612 in the PSID ("Value of all debt") . Homeowner x Debt is the interaction of the two. Skill is a dummy for whether the individual has received more than 12 years of education. Occ. tenure is years of uninterrupted tenure in broad occupation. Additional controls are: region dummies, age, age squared, unionization status, married status, year. Standard errors are clustered at the individual level. Columns with probit model show average marginal effect.

Table 11: Mapping of SCF occupation dummies into broad occupational groups.

Broad Group	Manual	Routine	Abstract
SCF occupation dummies	3	2, 4, 5, 6	1

or

$$m_{it} \leq \frac{y_{it}}{2} - \underline{m}_{it} \ ,$$

where $\underline{m}_{it}$ corresponds to a household's borrowing constraint, which, in line with Kaplan, Violante, et al. (2014)'s baseline definition, I assume to be one times monthly income.

B.5.2 Mapping of occupations

From 1989 until 2001, the SCF used the 3-digit 1980 and 1990 Census occupation codes, which were both very similar. The public-use files contain identifiers for whether the household head worked in one of six broad occupational groups. I map these six groups as closely as possible to the three groups used throughout this paper (manual, routine, abstract) using the consistent occupation-classification scheme proposed in the data appendix of Dorn (2009) (*Occ1990dd*).

Table Appendix 11 lists the assignment of SCF groups to the three broad occupational groups. While the overlap is very large and the assignment therefore unambiguous for the SCF groups 1 and 3 to 6, group 2 contains some abstract and some routine occupations (codes 203–389 of the 1980 Census classification). I therefore further look at the employment shares in each of these occupations provided by Autor and Dorn (2013) for the years 1990 and 2000. A significantly smaller fraction of workers in this group was employed in abstract occupations (8-9%, codes 203–258) than in routine occupations (21–22%, codes 274–389). Hence, I classify all workers of group 2 as routine, which also explains why I slightly overestimate the routine employment share in Figure 9 in the first years.

From 2004 onward, the SCF used the 2000 and 2010 Census codes, which were very similar to each other but quite distinct from the earlier Census codes. The SCF then assigned households again into six broad groups. While the overlap between the three broad occupational groups and the six SCF groups is somewhat weaker than prior to 2004, I still find that the assignment displayed in Table 11 yields the closest mapping.

To demonstrate that my classification of occupations is close to that of Autor and Dorn (2013), Figure 9 plots the self-constructed shares of employment obtained from the SCF next to those in Autor and Dorn (2013). While I am overestimating (underestimating) slightly the share of workers

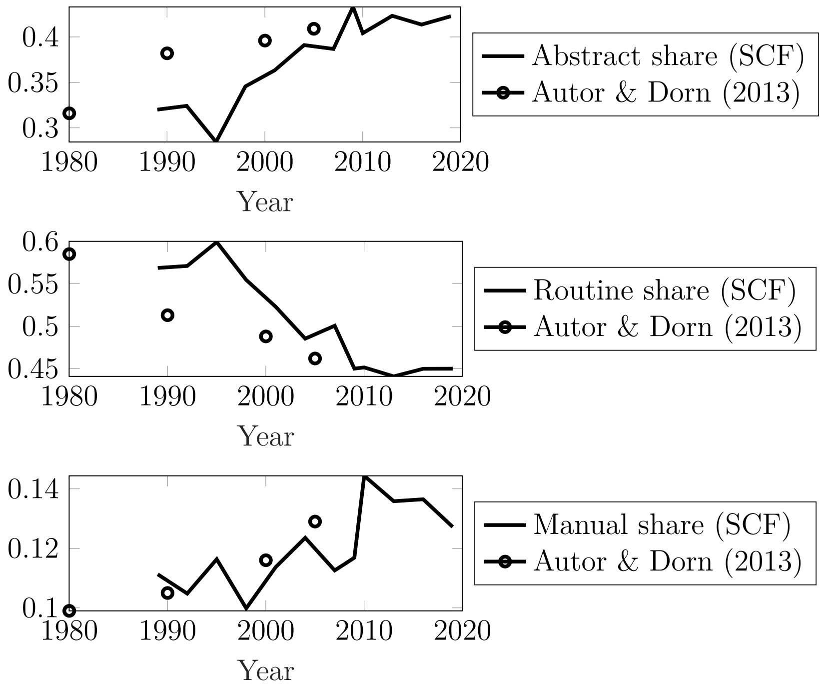

Figure 9: Employment shares in the SCF and in Autor and Dorn (2013).

employed in routine (abstract) occupations in the early 1990s, the time series are relatively closely aligned towards the early 2000s.

B.5.3 Time series of hand-to-mouth shares

Figure 10 plots the shares of households that I classify as being hand-to-mouth separately for each broad occupational group as well as unconditionally for all households. Two results stand out. First, there is a clear ordering of hand-to-mouth shares by occupational group, with workers in the manual occupations most likely to be hand-to-mouth, and abstract workers least likely. This is perhaps not surprising, given that abstract workers are usually the ones earning the highest incomes, manual workers earning the lowest incomes, and routine workers in between.

Second, routine workers, while not as likely as manual workers to be hand-to-mouth, are still more likely to be so than the average US household. Across all years, the probability of routine workers to be hand-to-mouth was on average 35%, higher than the average probability across all households of 29%. At each single point in time, the two series differ by three to seven percentage points.

Figures 11 and 12 further document that the high share of hand-to-mouth households among the routine workers is driven to a large extent by the high prevalence of wealthy hand-to-mouth households among them. These are households who are by definition hand-to-mouth, but own positive illiquid wealth such as equity in houses or indirect stock holdings (Kaplan, Violante, et al., 2014). In all years between 1995 and 2016, rou-

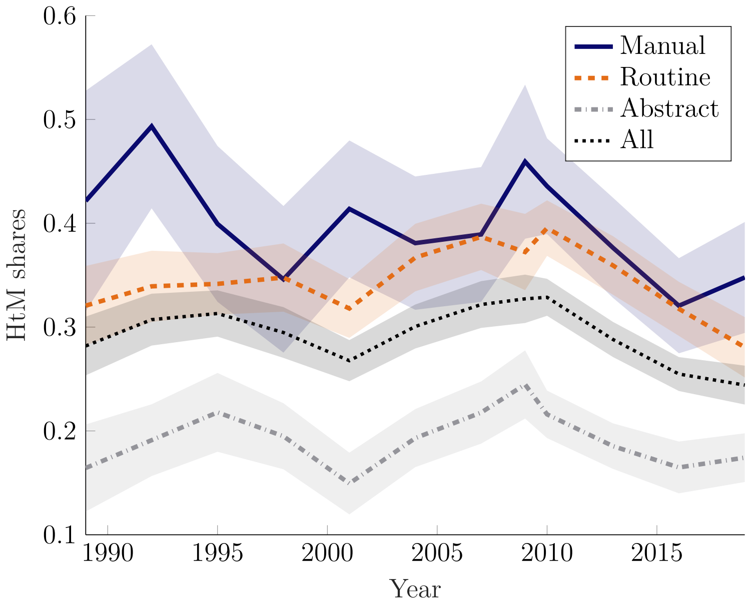

Figure 10: Hand-to-mouth shares in the three broad occupational groups.
Notes: Confidence intervals are at the 95% level. "All" refers to households whose head works in either of the three occupations, i.e. is employed.

tine workers were more likely to be wealthy hand-to-mouth than either abstract or manual workers.

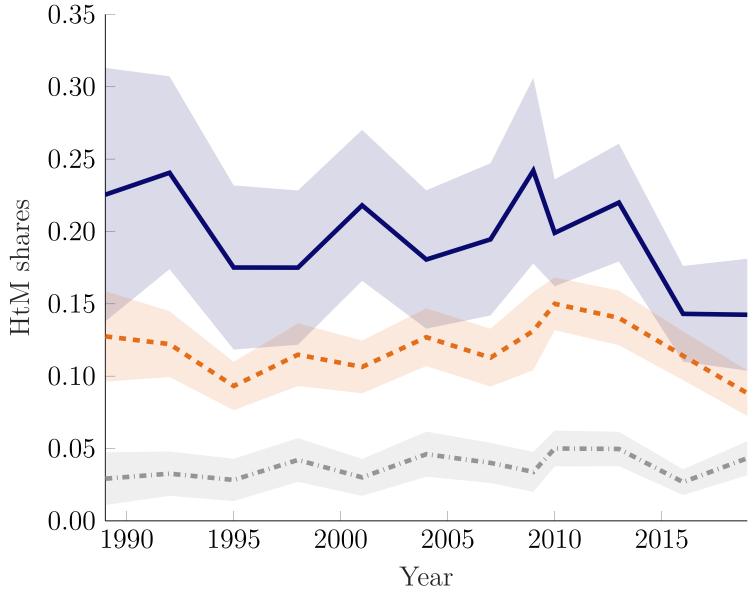

Figure 11: Poor hand-to-mouth shares in the three broad occupational groups.
Notes: Confidence intervals are at the 95% level.

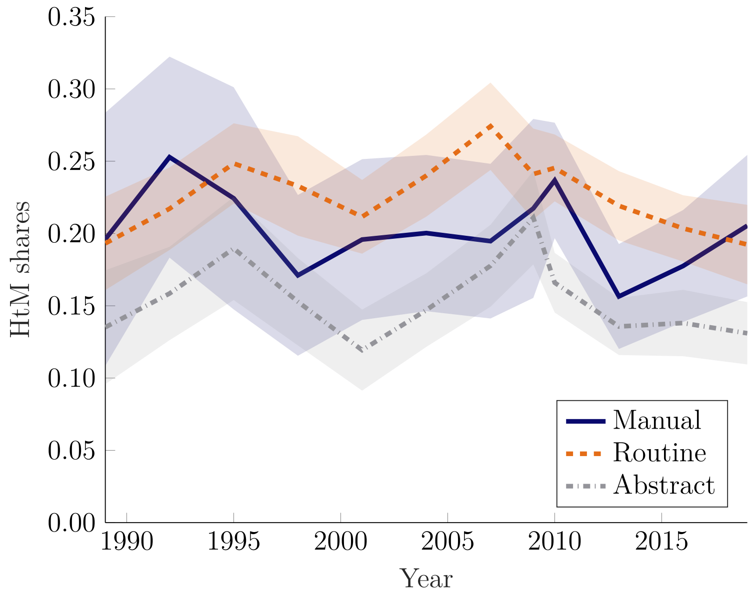

Figure 12: Wealthy hand-to-mouth shares in the three broad occupational groups.
Notes: Confidence intervals are at the 95% level.

C Details of the Full Model

C.1 FOCs of the representative firm

$$w_m = K_s^{\alpha}(1-\alpha)\tilde{Y}^{-\alpha}\mu_m\left(\frac{\tilde{Y}}{N_m}\right)^{\frac{1}{\gamma_m}} \tag{11}$$

$$w_r = \Omega(1-\mu_a)R^{\frac{\gamma_a-\gamma_r}{\gamma_a\gamma_r}}\mu_r N_r^{\frac{-1}{\gamma_r}} \tag{12}$$

$$w_a = \Omega\mu_a N_a^{\frac{-1}{\gamma_a}} \tag{13}$$

$$r_{ict} = \Omega(1-\mu_a)R^{\frac{\gamma_a-\gamma_r}{\gamma_a\gamma_r}}(1-\mu_r)K_{ict}^{\frac{-1}{\gamma_r}} \tag{14}$$

$$r_s = \alpha K_s^{\alpha-1}\tilde{Y}^{1-\alpha} \tag{15}$$

where

$$\Omega = K_s^{\alpha}(1-\alpha)\tilde{Y}^{-\alpha+\frac{1}{\gamma_m}}(1-\mu_m)\left(\mu_a N_a^{\frac{\gamma_a-1}{\gamma_a}} + (1-\mu_a)R^{\frac{\gamma_a-1}{\gamma_a}}\right)^{\frac{\gamma_m-\gamma_a}{(\gamma_a-1)\gamma_m}}$$

$$\tilde{Y} = \left[\mu_m N_m^{\frac{\gamma_m-1}{\gamma_m}} + (1-\mu_m)\left[\mu_a N_a^{\frac{\gamma_a-1}{\gamma_a}} + (1-\mu_a)R^{\frac{\gamma_a-1}{\gamma_a}}\right]^{\frac{\gamma_a(\gamma_m-1)}{(\gamma_a-1)\gamma_m}}\right]^{\frac{\gamma_m}{\gamma_m-1}}$$

C.2 Equilibrium

An equilibrium is defined as paths for household decisions $\{\tilde{k}_t, m_t, d_t, c_t, \ell_t, j_t\}_{t\geq 0}$, input prices $\{w_{m,t}, w_{r,t}, w_{a,t}, r_{ict,t}, r_{s,t}\}_{t\geq 0}$, government taxes and transfers $\{\tau_{l,t}, T_t\}_{t\geq 0}$, distributions $\{\Gamma_t\}_{t\geq 0}$, and aggregate quantities such that, at every t:

1. Given prices, aggregate quantities, the distribution $\{\Gamma_t\}_{t\geq 0}$, and the stochastic processes for individual states, policy functions c_t^*, ℓ_t^*,d_t^*, m_t^*, $\tilde{k}_t^*$ and j_t^* solve the households' problem (4).

2. The representative firm optimizes, given input prices. The FOCs (11), (12), (13), (14), and (15) hold.

3. The government budget constraint holds.

4. The labor markets clear, i.e. for $j \in \{m, r, a\}$:

$$N_{j,t} = \int_{i:j_t^*=j} y_{j,t}\ell_t^*\,\mathrm{d}\Gamma_t(i)$$

5. The no-arbitrage condition holds and the capital market clears:

$$\frac{K_{ict,t}}{q_{ict,t}} + K_{s,t} = \int_i \tilde{k}_t^*\,\mathrm{d}\Gamma_t(i)$$

6. The liquid asset market clears:

$$M_t^s = \int_i m_t^* \, \mathrm{d}\Gamma_t(i)$$

7. The resource constraint holds:

$$Y_t = C_t + I_{s,t} + I_{ict,t} + G_t + \int_i \chi(\cdot) + \kappa \max\{-m_t^*, 0\} \, \mathrm{d}\Gamma_t(i)$$

where C_t denotes aggregate consumption.

8. The sequence of distributions satisfies aggregate consistency conditions.

C.3 Hamilton-Jacobi-Bellman equation

For the specific processes of the exogenous productivity shocks used in the calibration, the solution to the problem of a household with state η_k is characterized by

$$\begin{aligned}
\hat{\rho} V_t(s, \eta_k, \epsilon, h, j, m, \tilde{k}) = & \max_{c,\ell,d} u(c, \ell) \\
& + V_{m,t}(s, \eta_k, \epsilon, h, j, m, \tilde{k}) \left[(1-\tau) w_j \ell y + \mathbf{1}_{m<0} \kappa m + T - d - \chi(d, \tilde{k}) - c\right] \\
& + V_{\tilde{k},t}(s, \eta_k, \epsilon, h, j, m, \tilde{k})(r\tilde{k} + d) \\
& + \lambda_\eta \left[V_t(s, \eta_{-k}, \epsilon, h, j, m, \tilde{k}) - V_t(s, \eta_k, \epsilon, h, j, m, \tilde{k})\right] \\
& + V_{\epsilon,t}(s, \eta_k, \epsilon, h, j, m, \tilde{k})(-\beta_\epsilon \epsilon) \\
& + \lambda_\epsilon \int_{-\infty}^{\infty} \left[V_t(s, \eta_k, x, h, j, m, \tilde{k}) - V_t(s, \eta_k, \epsilon, h, j, m, \tilde{k})\right] \phi(x) \, \mathrm{d}x \\
& + \dot{V}_t(s, \eta_k, \epsilon, h, j, m, \tilde{k}) \\
& \quad \text{such that:} \\
& \quad V_t(s, \eta_k, \epsilon, h, j, m, \tilde{k}) \geq \max_{\tilde{j} \in \{m,r,a\}} V_t(s, \eta_k, \epsilon, \underline{h}, \tilde{j}, m, \tilde{k})
\end{aligned}$$

and symmetrically for households with state η_{-k}. $\phi(\cdot)$ denotes the pdf of a normal distribution with standard deviation σ_ϵ. I employ the methods outlined in Achdou et al. (2022) to solve this household problem.[38]

C.4 Externally calibrated parameters

See Table 12.

[38] To deal with the stopping-time nature of the problem I further rely on the note "Liquid and Illiquid Assets with Fixed Adjustment Costs" by Greg Kaplan, Peter Maxted and Benjamin Moll, accessed on March 5, 2023, at https://benjaminmoll.com/wp-content/uploads/2020/06/liquid_illiquid_numerical.pdf.

Table 12: Externally calibrated parameters

Parameter	Value	Description	Source or Target
ζ	$1/(4 \cdot 45)$	death rate	avg. lifetime 45 years
φ	2.2	labor disutility	avg. labor time 8h/day
γ	1	elast. labor supply	Frisch elasticity of 1
$\underline{m}$	avg. qrtly. inc.	borr. constr.	Kaplan, Moll, et al. (2018)
τ_l	0.3	labor tax	Kaplan, Moll, et al. (2018)
T_t	$6\% \cdot Y_t$	lump-sum transfer	Kaplan, Moll, et al. (2018)
α	0.34	Non-ICT cap. share	Eden and Gaggl (2018)
δ_{ict}	0.175/4	deprec. ICT capital	Eden and Gaggl (2018)
δ_s	0.073/4	deprec. struct. capital	Eden and Gaggl (2018)
$[a_m, a_r, a_a]$	$[0, 0.18, 0.77]$	occ. prod. functions	vom Lehn (2020)
$F(s)$	$N(0,1)$	skill distrib.	vom Lehn (2020)

Notes: Rates are expressed as quarterly values.

C.5 Computation of counterfactual densities

To compute the counterfactual densities in Section 4 I proceed as follows. I start with the distribution of households over the state space in 1980 $\Gamma_{1980}(i)$. Denote the mass of workers of skill type $\tilde{s}$ in occupation j in 1980 as follows:

$$g_{1980,in,j,\tilde{s}} = \int_{i:\{h=\underline{h} \wedge s=\tilde{s}\}} \mathbf{1}_{\{j^*_{1980}(i)=j\}} \mathrm{d}\Gamma_{1980}(i)$$
$$g_{1980,ex,j,\tilde{s}} = \int_{i:\{h=\bar{h}_j \wedge s=\tilde{s}\}} \mathbf{1}_{\{j^*_{1980}(i)=j\}} \mathrm{d}\Gamma_{1980}(i)$$

where $j^*_{1980}(i)$ denotes optimal occupational choices in 1980, and in and ex in the index of the densities abbreviates "inexperienced" and "experienced", respectively.

Note that when solving for the transition path I discretize time into periods (Kaplan, Moll, et al., 2018). For each point t on the discretized time grid I compute for each skill type $\tilde{s}$ the average probability across all inexperienced households to choose a certain occupation j, and denote this probability by $p_{t,\tilde{s},j}$, i.e.

$$p_{t,\tilde{s},j} = \frac{\int_{i:\{h=\underline{h} \wedge s=\tilde{s}\}} \mathbf{1}_{\{j^*_t=j\}} \mathrm{d}\Gamma_t(i)}{\int_{i:\{h=\underline{h} \wedge s=\tilde{s}\}} \mathrm{d}\Gamma_t(i)} \tag{16}$$

This effectively averages over the wealth (m and $\tilde{k}$) and idiosyncratic productivity (η and ϵ) dimensions, while conditioning on human capital and skill type. Γ_t refers to the distribution at time t during the baseline transition. Similarly, for $j \in \{m, r, a\}$, I compute the average probability $x_{t,\tilde{s},j}$ that experienced households exit their current occupation j and become

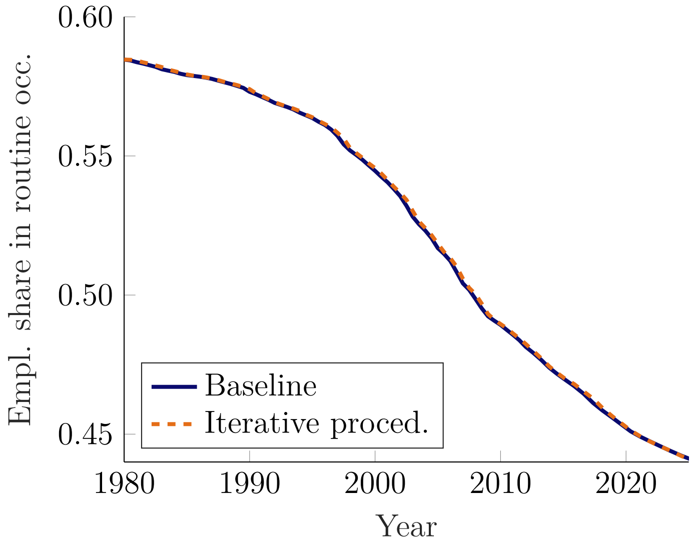

Figure 13: Employment share in the routine occupation.

Notes: Blue solid: true employment share during the transition ($\int_{i:j_t^*=r} \mathrm{d}\Gamma_t(i)$). Orange dashed: employment share constructed from iterative procedure ($\sum_s g_{t,in,r,s} + g_{t,ex,r,s}$).

inexperienced households

$$x_{t,\tilde{s},j} = \frac{\int_{i:\{j_{t-1}^*=j \wedge h=\bar{h}_j \wedge s=\tilde{s}\}} \mathbf{1}_{\{j_t^* \neq j\}} \mathrm{d}\Gamma_t(i)}{\int_{i:\{j_{t-1}^*=j \wedge h=\bar{h}_j \wedge s=\tilde{s}\}} \mathrm{d}\Gamma_t(i)} \tag{17}$$

I then use these probabilities to iterate forward the densities of skill type $\tilde{s}$ in the following way. For inexperienced households in occupation j:

$$\begin{aligned} g_{t+1,in,j,\tilde{s}} = p_{t,\tilde{s},j} \Bigg\{ & \underbrace{\sum_{k \in \{m,r,a\}} \mathrm{e}^{-\lambda_h^k \cdot \mathrm{dt}} g_{t,in,k,\tilde{s}}}_{\text{inexp. HHs who do not become experienced}} \\ & \underbrace{+ \sum_{k \in \{m,r,a\}} \left[x_{t,\tilde{s},k} \left(\mathrm{e}^{-\zeta \cdot \mathrm{dt}} g_{t,ex,k,\tilde{s}} + (1 - \mathrm{e}^{-\lambda_h^k \cdot \mathrm{dt}}) g_{t,in,k,\tilde{s}} \right) \right]}_{\text{exp. HHs who exit their occ.}} \\ & \underbrace{+ (1 - \mathrm{e}^{-\zeta \cdot \mathrm{dt}}) \sum_{k \in \{m,r,a\}} g_{t,ex,k,\tilde{s}}}_{\text{dead exp. HHs}} \Bigg\} \end{aligned} \tag{18}$$

And for experienced households in occupation j:

$$g_{t+1,ex,j,\tilde{s}} = (1 - x_{t,\tilde{s},j}) \left\{ \mathrm{e}^{-\zeta \cdot \mathrm{dt}} g_{t,ex,j,\tilde{s}} + (1 - \mathrm{e}^{-\lambda_h^j \cdot \mathrm{dt}}) g_{t,in,j,\tilde{s}} \right\} \tag{19}$$

The last step is to aggregate these densities across all skill types s.

Figure 13 plots both the actual mass of households in the routine occupations (i.e. using Γ_t), as well as the one obtained from the iterative procedure described here. It reveals that by using the iterative procedure I recover the actual employment share during the transition very well.

Figure 14 For this counterfactual, in which I assume that only newborns can make occupational choices, I first set all exit probabilities $x_{t,s,j} = 0$ in (18) and (19). Furthermore, only newborns can choose a new occupation, i.e. (18) becomes

$$g_{t+1,in,j,\tilde{s}} = \mathrm{e}^{-\zeta \cdot \mathrm{d}t} \mathrm{e}^{-\lambda_h^j \cdot \mathrm{d}t} g_{t,in,j,\tilde{s}} + p_{t,\tilde{s},j}(1 - \mathrm{e}^{-\zeta \cdot \mathrm{d}t}) \sum_{k \in \{m,r,a\}} \left[g_{t,ex,k,\tilde{s}} + \mathrm{e}^{-\lambda_h^k \cdot \mathrm{d}t} g_{t,in,k,\tilde{s}} \right] .$$

Figure 4 To arrive at the counterfactual employment shares that use policy functions of the liquid-wealthy households only, I replace $\Gamma_t(i)$ in (16) and (17) with

$$\Gamma_t(i) \cdot \mathbf{1}_{i:\{m \geq 70.(90.) \text{ prctl.}\}} .$$

Households with high liquid assets have systematically different (i.e. higher) productivity draws (ϵ and η) than the average population. I correct for this as follows: In each t, I compute (16) and (17) for each point on the discretized grids of the two productivity processes. I then weight these probabilities according to the invariant stationary distribution of ϵ and η in the population. This way, I arrive at probabilities $p_{t,\tilde{s},j}$ and $x_{t,\tilde{s},j}$ that correspond to a household of average productivity but high liquid wealth. I then iterate on (18) and (19) as described above.

D Untargeted moments

Hand-to-mouth shares by occupation The calibration targets unconditional shares of hand-to-mouth households in the economy, but not the shares conditional on occupation. These moments are reported in Table 13 for the initial steady state of the model as well as for the earliest available wave of the SCF (1989). Overall, the model provides a very good fit to the data in this regard.

Table 13: Hand-to-mouth shares by occupation.

	SCF (1989)			Model (init. steady state)		
	HtM	wealthy HtM	poor HtM	HtM	wealthy HtM	poor HtM
Manual	42%	20%	22%	41%	26%	15%
Routine	32%	20%	13%	34%	23%	11%
Abstract	16%	13%	3%	17%	12%	6%

Notes: Wealthy HtM are hand-to-mouth households with positive illiquid assets, poor HtM are those with zero illiquid assets.

Table 14: Switching decision and liquid asset holdings (model).

	(1)	(2)	(3)	(4)	(5)	(6)
HtM	-0.0063***	-0.0076***	-0.0060***	-0.0084***		
	(0.0013)	(0.0013)	(0.0013)	(0.0012)		
Occupational tenure	-0.055***	-0.048***	-0.039***	-0.033***	-0.039***	-0.033***
	(0.0032)	(0.0030)	(0.0019)	(0.0016)	(0.0019)	(0.0017)
Skill		0.068***		0.064***		0.064***
		(0.0020)		(0.0020)		(0.0020)
Poor HtM					-0.0034	-0.0091***
					(0.0028)	(0.0025)
Wealthy HtM					-0.0066***	-0.0082***
					(0.0015)	(0.0013)
Time trend	Yes	Yes	Yes	Yes	Yes	Yes
Observations	99491	99491	99491	99491	99491	99491
Model	OLS	OLS	Probit	Probit	Probit	Probit

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Data are from a synthetic panel of 10,000 households, simulated from the model at annual frequency between 1980 and 2020. Dependent variable is whether or not individual leaves routine occ. between t and $t+2$. Occ. tenure is a dummy for whether or not worker is experienced. Skill is s. I only use observations in or after 1999. Standard errors are clustered at the individual level. Columns with probit model show average marginal effect.

Regressions of switching decision on hand-to-mouth status Adding to Section 5, another way to show that individual changes in behavior due to borrowing constraints become visible at the aggregate level, is to rerun the regressions from the end of Section 3 (equation (1)). This time, instead of using data from the PSID, I use the synthetic panel of 10,000 households described in the calibration. Mimicking the empirical implementation, I regress a dummy for whether a household switches between year t and $t+2$ on a dummy for whether a household is hand-to-mouth. I use the same controls as in the empirical specification (as far as they exist in the model), and use observations in or after 1999. Table 14 shows the results. The same negative association between current liquid asset holdings and future switching decision found in the data also prevails in the model.[39] Also as in the data, the correlation is negative both for poor and for wealthy hand-to-mouth households.

The role of newborns in the decline of the routine employment share Cortes, Jaimovich, Nekarda, et al. (2020) provide empirical evidence that 34%–43% of the fall in routine employment has been due to a

[39] The point estimates are smaller (in absolute terms) than their empirical counterparts in Table 2. This owes to the somewhat lower level of occupational mobility in the model compared to the PSID (see discussion below): as the overall level of mobility is lower in the model, so are the marginal effects. However, despite the marginal effects being closer to zero, all of them are statistically significant.

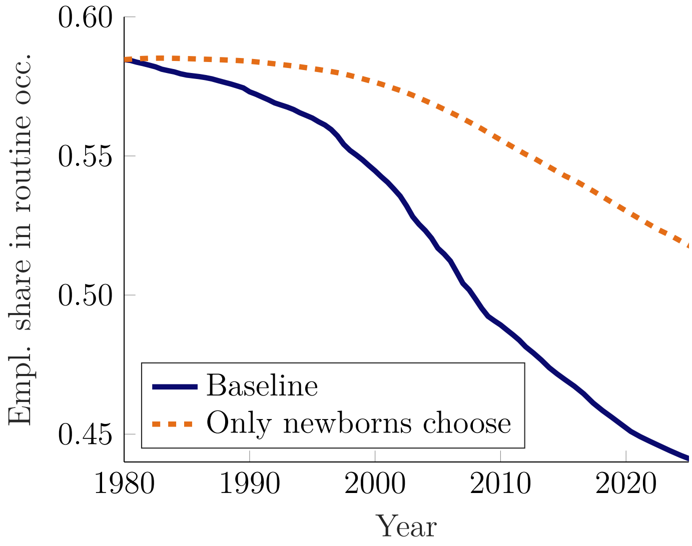

Figure 14: Employment share in the routine occupation
Notes: Blue solid: baseline transition. Orange dashed: counterfactual with initial occupational choices made permanent for rest of life.

reduced propensity, especially by young workers, to enter routine occupations from non-employment and unemployment.[40] While I do not model unemployment, one can interpret the entry of newborn households as entry from non-employment. I therefore conduct a counterfactual exercise to assess whether my model performs realistically in this regard.

Starting in the initial steady state, the optimal occupational policy functions j_t^* dictate which occupation new labor market entrants choose at the beginning of their life. Afterwards, workers never switch the occupation, i.e. I iterate forward the distribution of households over the state space assuming that households only make an occupational choice at the beginning of their lives. Hence, this is an out-of-equilibrium exercise that abstracts from general equilibrium effects on initial occupational choices and on wages, just as in Cortes, Jaimovich, Nekarda, et al. (2020). Details on how I construct this and all following counterfactuals are relegated to Appendix C.5.

The orange dashed line in Figure 14 depicts the resulting mass of routine workers in the economy. It falls much more slowly compared to the baseline transition (blue solid), which implies that some of the fall in routine labor in the model is due to net occupational switching. Consistent with Cortes, Jaimovich, Nekarda, et al. (2020), in 2020 41% of the fall in routine employment since 1980 is accounted for in the model by the re-

[40] They find that outflows to non- and unemployment are not important for explaining the decline. Also, the authors note that 54–60% of the drop in routine employment cannot be explained by changed exit to and entry from non- and unemployment, and is therefore due to job-to-job transitions and other transitions not covered by their decomposition. In my model, job-to-job transitions and reduced inflows from new labor market entrants are the two only drivers of the decline in routine employment.

duced propensities of new labor market participants to enter the routine occupation.

Timing of the fall in routine employment The magnitude of the decline of routine employment from 58.5% in 1980 to 45.0% in 2020 was targeted in the calibration. The dynamics of the fall, though untargeted, are however well in line with empirical evidence in Cortes, Jaimovich, and Siu (2017), Jaimovich and Siu (2020), and Kikuchi and Kitao (2020): Routine employment declines gradually prior to the 2000s and faster afterwards.

Occupational mobility Mobility across the three broad occupations averages 1.0% annually along the transition in the model. This is lower than mobility in the PSID, which between 1979 and 1997 was 5.0% for males. However, it is well-known that due to measurement errors estimates of occupational mobility in the PSID are inflated (Kambourov and Manovskii, 2008). Using CPS data, Kikuchi and Kitao (2020) estimate occupational mobility of 1.3% for abstract, 1.4% for routine and 5.6% for manual workers, closer to the average value in my model.[41]

Relative wages and interest rate over time While I target the change in employment shares in the calibration, the evolution over time of relative wages per efficiency unit is untargeted. Cortes (2016) estimates that the log of w_a/w_r rose by about 25% between 1980 and 2007, and the log of w_m/w_r by 10 to 15%.[42] The model produces a nearly identical rise in the log of w_a/w_r (23%), though a somewhat smaller increase in the log of w_m/w_r (4%). This indicates that the model captures very well the incentives over time for relatively high-skilled routine workers to switch to the abstract occupation. The incentives for low-skilled routine workers to switch to the manual occupation is a bit smaller in the model than in the data. If anything, this could lead me to underestimate the benefits of policies targeted at switchers to the manual occupation.

Figure 15 further plots the evolution of the interest rate. Since firms demand more capital when its relative price falls, the interest rate rises over time. This is in line with empirical evidence discussed in Moll et al. (2022).

[41]These numbers correspond to the high-skilled group in the case of abstract and to the low-skilled group in the case of routine and manual workers reported in Kikuchi and Kitao (2020). These groups respectively account for the largest portion of employment in each occupation.

[42]These numbers correspond to the lower left panel of Cortes (2016, Figure 6), as in that specification he allows for occupation-specific tenure profiles, as I do in my model.

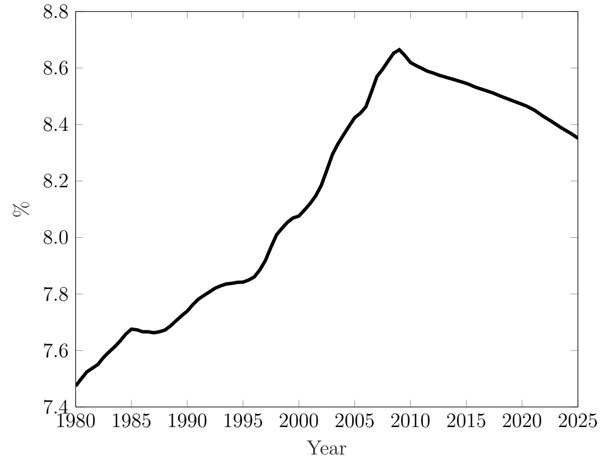

Figure 15: Model-implied (annual) interest rate r (in %).

Wage changes upon occupational switch In the calibration I target the average wage change of workers leaving the routine for either the manual or the abstract occupation, compared to workers who stayed in the routine occupation. Remembering Table 1, Cortes (2016) also compares wages over longer horizons than one year, documenting faster wage growth for switchers than for stayers. I visualize his estimates, i.e. the values from Table 1, and the corresponding statistics from the simulated panel of my model in Figure 7. Overall, I am able to capture the differential wage paths of switchers compared to stayers quite well, though for both directions of the switch long-run wage growth is a bit less pronounced in the model than in the data. In the case of switches to the manual occupation, this owes in part to the somewhat smaller than empirically observed increase in the log of w_m/w_r mentioned in the previous paragraph.

E Details on Policies Targeted at Switchers to Abstract Occupation

E.1 Finding the constraint optimal policies

The blue solid lines in Figure 16 show consumption-equivalent welfare changes according to (5) averaged over time for different values of R (left panel) and L (right panel). In both cases, welfare differences are inversely u-shaped in the size of the policy, with the optimal monthly wage replacement being 420$ and the optimal loan being 10,000$.[43] Furthermore, the optimal wage replacement leads to a larger increase in welfare than the

[43]Conditional on observing a year-on-year earnings loss when switching from the routine to the abstract occupation, the median monthly earnings loss in the synthetic panel simulated from the model is 680$. Hence, the optimal wage replacement program replaces about 60% of the median earnings loss.

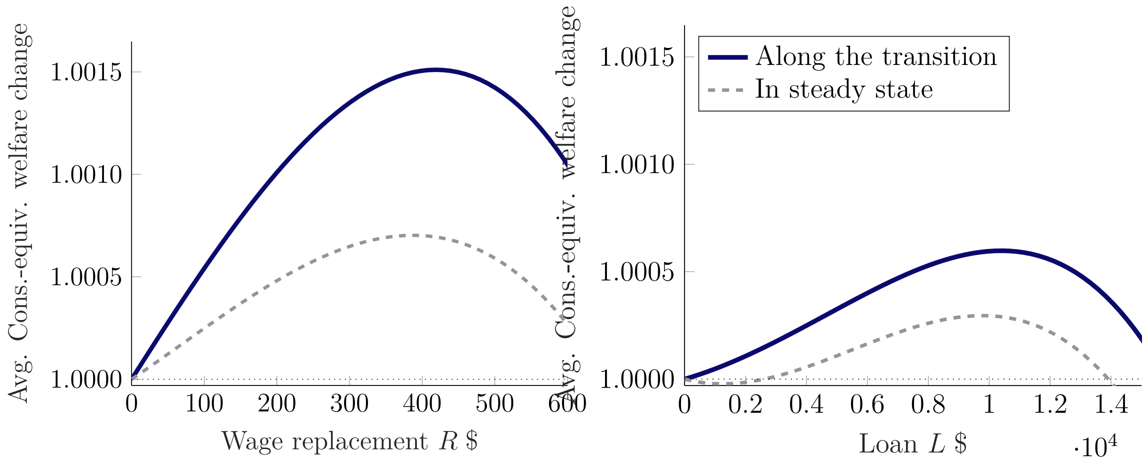

Figure 16: Average change in welfare under different policies compared to the baseline

Notes: Cons.-equiv. welfare change ϕ_t (defined in equation (5)), averaged between $t = 1980$ and $t = 2070$. "Transition": baseline transition path including technological change ($q_{ict,t}$ rises). "Steady state": transition during which no technological change occurs ($q_{ict,t}$ stays constant at its 1980 level). Nominal quantities are expressed in 2017 US dollars.

optimal loan. The former achieves an average increase in consumption of 0.15%. This masks significant variation across cohorts and especially across skills, though, as shown in Section 6.2.3.

E.2 Output, ICT capital, and labor productivity

Figure 17 shows the effect of the constrained optimal policies on output, ICT capital and labor productivity. The left panel plots relative deviations of output Y_t under the policies from its values in the baseline transition. During the first years, especially under the wage replacement program, effects on output are slightly negative. This owes to higher labor taxes and additional switching of experienced workers which destroys human capital in the short term. After 2000, both policies have positive effects on output which persist far into the future. Averaged over the years 1980 to 2050, both policies cause output to rise by about 0.06% compared to the baseline, with a peak of about 0.15% during the later years of the transition.[44] This is in line with the results from the tractable model in Appendix A.2. Alleviating the borrowing constraints leads to a more efficient allocation of labor across occupations and therefore benefits output.

The middle panel shows that the constrained optimal policies lead to a crowding in of ICT capital of up to 1% (3%). As there is less supply of routine work, which is a substitute to ICT capital, and more supply of abstract work, which is a complement, firms optimally react by using more

[44] The dip in 2025 is explained by the fact that the programs are phased out in that year (see design of the policies). This causes some last workers to take advantage of the policies, leading to a spike in labor taxes to finance this (see Figure 28).

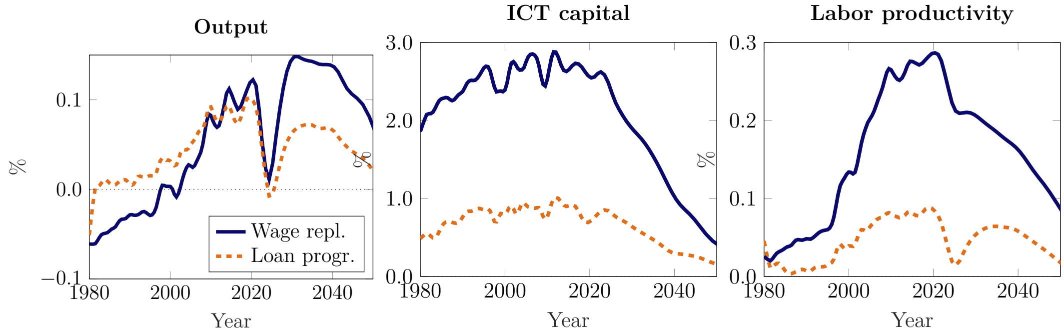

Figure 17: Deviations of variables under constrained optimal policies relative to transition without policy

Notes: Left: Output Y_t. Middle: ICT capital $K_{ict,t}$. Right: Labor prod. $Y_t / \left(\int_i \ell_t^* \, \mathrm{d}\Gamma_t(i)\right)$.

ICT capital in production when the policies are in place. Finally, as the right panel shows, the policies have a positive impact on labor productivity, as measured by output divided by hours worked. This is driven by the increase in ICT capital as well as the improved allocation of workers across occupations, i.e. more abstract and less routine labor. These two factors are intimately linked, given the complementarities inherent to the aggregate production function (2): the reallocation of labor towards the abstract occupation only leads to productivity gains if accompanied by an increase in ICT capital.[45]

E.3 Implementing the policies in the steady state

The grey dashed lines in Figure 16 show the welfare implications of introducing the policies in the initial steady state of the model.[46] About half of the welfare gains obtained along the actual transition path are also realized in the steady state. There are two main reasons for this. First, the policies provide insurance to eligible workers and redistribute across households just like they do when implemented along the transition. Second, there can be some efficiency gains even in steady state: while wages are constant in steady state, returns to tenure $(\bar{h}_j - \underline{h})$ are higher and hence earnings paths steeper in the abstract than in the routine occupation. Therefore, even in steady state the borrowing constraint prevents some workers from optimally borrowing against high future income and choosing the abstract instead of the routine occupation.

[45] When fixing the capital stock at its values under the baseline transition and only letting labor inputs $N_{j,t}$ adjust to their new values under the policies, the effect on output is negative (not shown). Only in conjunction with the increase in ICT capital do the policies cause output gains.

[46] Technically, I still simulate a transition to compute the welfare criterion, but one in which q_{ict} stays constant at its 1980 level. I then phase out the policies between 2025 and 2070, which makes the results comparable to the policy analysis along the actual transition path with a falling relative price of ICT capital.

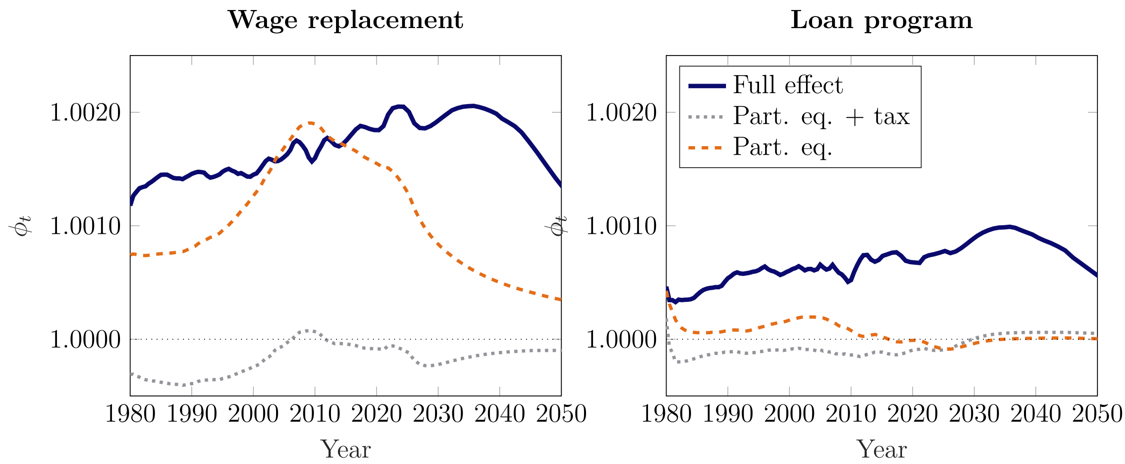

Figure 18: Consumption-equivalent welfare change under constrained optimal policies compared to baseline over time

Notes: Cons.-equiv. welfare change ϕ_t (defined in equation (5)). Left: optimal wage replacement. Right: optimal loan program. "Part. eq." holds fixed taxes $\tau_{l,t}$ and factor prices at their values from the baseline transition without policy. "Part. eq. + tax" holds fixed factor prices only. "Full effect" is the full general equilibrium effect.

In sum, however, the welfare gains are only half of those obtained when implementing the policies along the transition path. In addition, unlike along the transition, the policies' effect on output is negative in steady state (graph not shown). This points to an important role of efficiency gains causing the welfare improvements along the transition.

E.4 Welfare over time and dissection of channels

How are the welfare gains from implementing the optimal wage replacement (loan) program along the transition distributed over time? To answer this question, Figure 18 plots the change in welfare in terms of consumption-equivalent units from equation (5). The blue solid line shows that both programs lead to welfare improvements along the entire transition path.

To shed further light on the channels that can cause welfare improvements discussed in the main text, I further dissect the total general equilibrium effect of the policy. The orange dashed line in Figure 18 plots ϕ_t when keeping the labor tax as well as all factor prices, i.e. wages and the interest rate, constant at their values from the baseline transition without policy. The only change compared to the baseline is that the policies are implemented (without being funded). As can be seen, this partial equilibrium effect raises welfare at all (most) horizons under the wage replacement (loan) program. Households benefit from the wage replacement payments (loans) and increase consumption and leisure. Only at later stages of the loan program, when no new households become eligible and earlier recipients have to pay back the loan they received, the welfare effect turns slightly negative.

On top of this partial equilibrium effect, the grey dotted line further lets the labor income tax $\tau_{l,t}$ adjust to the level needed to finance the payments. This takes away from the welfare gains obtained initially, and the overall effect becomes slightly negative at most horizons. This clearly demonstrates that the higher insurance as well as the redistribution channel are not the key drivers of the aggregate welfare gains obtained from implementing the constrained optimal policies. The blue solid line plots the full general equilibrium effect, i.e. after wages and the interest rate have adjusted to clear the factor markets. This effect is positive and quantitatively the most important one at almost all horizons of both programs.

F Robustness Checks

F.1 Preference heterogeneity

In the baseline model all households share the same preferences. To obtain heterogeneity in asset holdings, households are exposed to uninsurable, idiosyncratic labor income risk (Aiyagari, 1994; Bewley, 1983; Huggett, 1993). Hence, if a household is close to the borrowing constraint (hand-to-mouth), it is because of "bad luck", i.e. low productivity draws.

In contrast, Aguiar et al. (2020) have recently argued that being hand-to-mouth is driven rather by preference heterogeneity, i.e. that there exist certain (preference) types of households who are especially likely to be hand-to-mouth. Using data from the PSID, they estimate three groups of preference types in the population, letting the groups differ by their discount factor and intertemporal elasticity of substitution. While the first two groups' impatience and elasticity of substitution are estimated at relatively standard values, the third group features a quite high impatience and elasticity of substitution. It is the latter group, which constitutes about a fifth of the population, that according to Aguiar et al. (2020) is most important for explaining the empirical share of poor hand-to-mouth agents in the economy.

Differences in preferences, in particular in the discount rate, would not subtract from the relevance of the mechanism that I propose in this paper. The mechanism put forward concerns optimal *investment decisions*, in the form of occupational choices, which depend on interest rates and wage paths, but not on preferences. Put differently, efficient occupational choices are determined independently of preferences. I show this formally in the tractable model in Section 2, see, for instance, Proposition 1, which does not depend on β. Intuitively, even the most impatient household would want to maximize lifetime income by, first, making the efficient occupational

Table 15: The three preference type groups, taken from Aguiar et al. (2020).

Group	Population share	Discount rate	Intert. elast. of subst.
I	44.7%	$\hat{\rho} = -\ln(0.97)/4 = 0.0076$	$1/\sigma = 0.53$
II	33.7%	$\hat{\rho} = -\ln(0.94)/4 = 0.0155$	$1/\sigma = 0.95$
III	21.6%	$\hat{\rho} = -\ln(0.72)/4 = 0.0821$	$1/\sigma = 2.87$

Notes: Aguiar et al. (2020)'s model is at annual frequency. They define the intertemporal elasticity of substitution as σ, while it is $1/\sigma$ in my model. For comparison, the benchmark model of Section 4 (without preference heterogeneity) has $\hat{\rho} = 0.0176$ and $1/\sigma = 1$.

choices over time, and second, consuming all lifetime income right away by borrowing against future income.[47] It is *being borrowing-constrained*, or hand-to-mouth, that makes this strategy suboptimal from the viewpoint of the household, inefficiently distorting occupational choices.

However, whether households are at the borrowing constraint because of low productivity draws or because of their preferences could still matter for the quantitative results of the model, in particular for the policy analysis. Therefore, I re-run the policy analysis, only this time with three fixed preference types in the population. I use the same calibration as in the main text and run the model once without any policy and once with the optimal wage replacement of 420$ found in Section 6. The three groups are taken from Aguiar et al. (2020). The implied parameters in the context of my model are listed in Table 15.[48]

Figure 19 shows that preference heterogeneity only matters very little for the welfare changes that the policy brings about. For each preference type, welfare changes in terms of expected lifetime consumption follow a very similar pattern across skill type as in the baseline model, both qualitatively and quantitatively.

While Group III households tend to experience the least pronounced welfare changes, this is not because workers of this preference type do not enter into the program. To the contrary, as Figure 20 shows, Group III households are in fact more likely than households of the other two groups to become eligible to the program by leaving the routine occupation as an experienced worker. This is intuitive, as Group III households are the most impatient ones, and hence they value the instant transfers paid under the program relatively highly.

Instead, that welfare changes are smaller for the impatient households is due to the fact that welfare is measured here as expected lifetime consumption at the beginning of a household's life as an *inexperienced* worker.

[47] Impatient households discount future *consumption*, not future *income*.

[48] Aguiar et al. (2020) use Epstein-Zin preferences to distinguish the intertemporal elasticity of substitution and risk aversion. Since I use a simpler utility function, risk aversion σ varies across preference groups in my case, depending on the calibrated intertemporal elasticity of substitution.

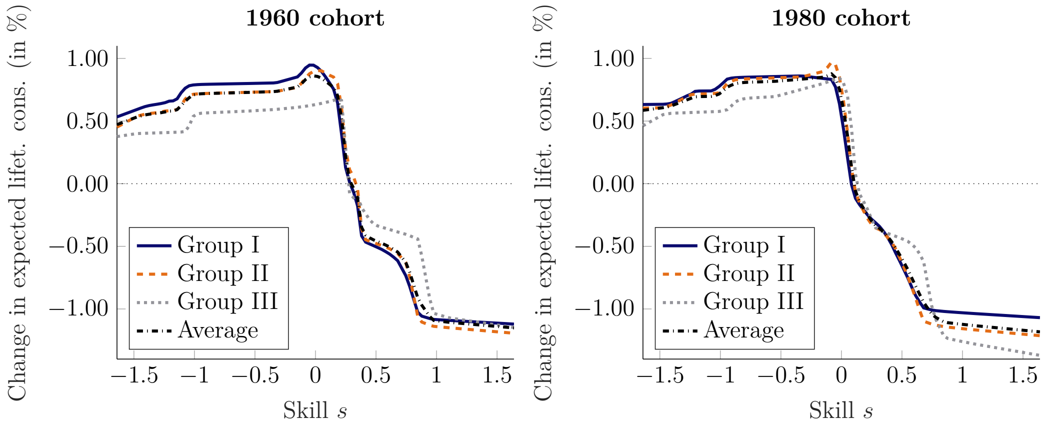

Figure 19: Expected lifetime consumption-equivalent welfare change of newborn households under wage replacement of 420$

Notes: $(\bar{\phi}_{\tilde{T}}(s) - 1) \cdot 100\%$ is on the y-axis, where $\bar{\phi}_{\tilde{T}}(s)$ is as defined in equation (6). "Average" is the average change in consumption equivalents, weighted by population share of the preference types.

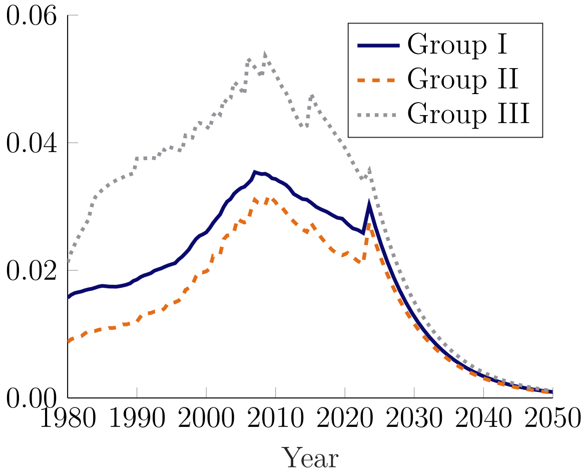

Figure 20: Share of eligible households within each preference type.

Hence, potential direct benefits from entering into the program as an *experienced* routine worker leaving the occupation occur in the relatively distant future, which impatient households care relatively little about.[49]

The fact that all three groups' welfare changes in a similar way reinforces the finding that it is the general equilibrium (or efficiency) gains caused by the policy, rather than the direct effects of program participation, that are the main driver of the welfare changes (see discussion in Section 6 and Appendix E). These efficiency gains, reflected on average in higher income, affect households regardless of their preferences.

Lastly, the wage replacement policy raises output even slightly more in the economy with preference heterogeneity than in the baseline economy with homogeneous preferences (graph not shown). This is another indication that the inefficiency highlighted in this paper is robust to the inclusion of a realistic degree of preference heterogeneity.

F.2 Myopia

In the baseline model, agents have perfect foresight about the path of q_{ict} once it begins its rise to a new, higher level. Analogously, they anticipate perfectly how prices (wages and the interest rate) are going to evolve. In its recurrent Occupational Outlook Handbooks the US Bureau of Labor Statistics does publish projections on how it expects employment to grow in hundreds of occupations and how they are affected by technological growth (see Edin et al., 2023, for a detailed discussion). Still, the assumption of perfect foresight about technological growth and how it affects wages is surely extreme.

Therefore, I conduct an experiment in which I assume instead that households are entirely myopic. In particular, at each point in time, they believe that the level of technology $q_{ict,t}$ and all prices will stay at their current levels forever.[50] In the next moment, households are then surprised by a further rise in q_{ict} and an according adjustment of prices. This implies that households do not foresee the fall in the routine relative to the manual and abstract wage at the early stages of the transition. I view this experiment as the opposite extreme case compared to perfect foresight.

As would be expected, the fall in the share of workers in the routine occupation is initially protracted under myopia compared to the perfect foresight case. Not foreseeing the continuing routine-biased technological

[49] Moreover, as households born in a given year but with different discount rates value utility at a certain point in the future differently, the way the policy alters the equilibrium wage paths (Figure 5) has a heterogeneous impact on different preference types' lifetime utility. This further explains the differences in welfare changes across preference types.

[50] Technically, prices at each t are computed such that all factor markets clear, given that households assume that these market-clearing prices stay in place forever.

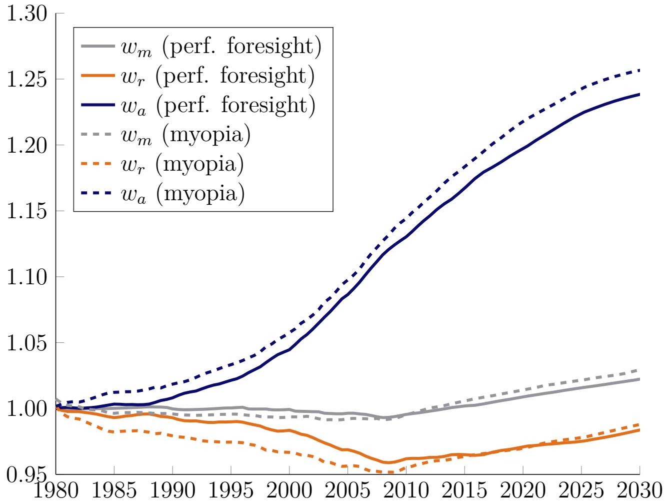

Figure 21: Wages per efficiency unit of labor

Notes: Solid lines: Baseline model (perfect foresight). Dashed lines: Myopia. Wages in occupation j are relative to w_j in 1980 in the baseline transition.

change, fewer workers find it optimal to leave the routine occupation. Accordingly, routine wages are depressed relative to the perfect foresight case at the beginning of the transition, as Figure 21 visualizes. This pattern reverses only at a later stage of the transition, when technology q_{ict} approaches its new higher steady state level. Not foreseeing the eventual rebound in routine wages, ever more workers leave the routine occupation, thereby putting upward pressure on the routine wage.

Since the supply of labor in the routine occupation is initially even larger under myopia than under perfect foresight (where it is already inefficiently large along the transition), implementing the wage replacement policy from Section 6 (420$) brings about even higher gains. First, average consumption increases by 0.21% under myopia, compared to 0.15% under perfect foresight (computed according to equation (5)). Second, the average output gain is more than twice as large as under perfect foresight. Figure 22 illustrates this, plotting output changes due to the policy, once under perfect foresight (blue solid) and once under myopia (orange dashed). This demonstrates that the assumption of perfect foresight likely understates the benefits of the proposed policies.

F.3 Further discussion

To keep the model computationally tractable, I made a number of simplifying assumptions. For instance, I have assumed that all occupation-specific human capital is irrevocably lost upon changing the occupation. Also, the rate at which workers gain new human capital, λ_j^h, is exogenous. Hence, there is no endogenous effort margin, by which households could influence the speed of human capital accumulation. Loosening these assumptions

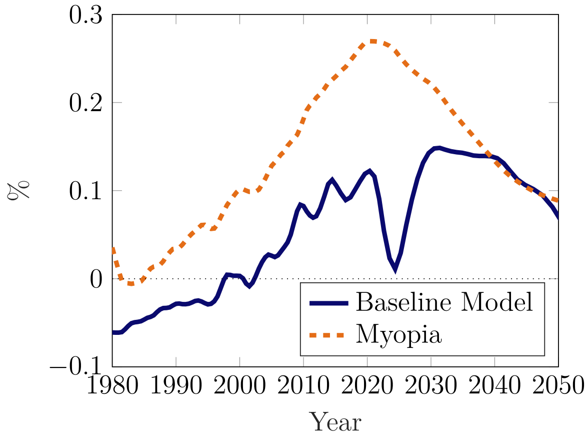

Figure 22: Deviation of Output Y_t under wage replacement of 420$ relative to transition without policy.
Notes: Blue solid: Baseline transition (perfect foresight). Orange dashed: Myopia.

would not affect the principle of the mechanism proposed in this article, i.e. when at the borrowing constraint, workers prefer to work in the occupation in which they can currently earn the highest wages. However, the quantitative results could change. For instance, the distortionary effect of the loan program could grow under partial recall of human capital: If workers could switch the occupation while only slowly losing old human capital, entering the program and returning to the routine occupation at a later point would become (even) more attractive. The wage replacement program would likely be less affected by partial recall because the transfer payments are conditioned on staying in the new occupation. This reinforces the point made in the main text that the wage replacement program is the more efficient of the two modelled programs. While it would be interesting to extend the model further along these dimensions, this lies outside the scope of the current article and is left for future work.

Moreover, I have followed Kaplan, Moll, et al. (2018)'s calibration for a subset of important parameters. For instance, as them, I set the Frisch elasticity of labor supply γ, as well as the intertemporal elasticity of substitution $1/\sigma$ equal to 1. Both values are common in the literature. However, the question remains how sensitive the results are to this particular parameter choice.

I therefore redo the analysis, once with a lower elasticity of $\gamma = 1/2$, and once with a lower elasticity of substitution $1/\sigma = 1/2$, both also values that are commonly found in the macroeconomics literature. All other parameters are as in the benchmark model. Figure 23 shows that wage paths, in particular their dynamics, are very similar across the three different specifications. Perhaps even more importantly, Figure 24 shows that the welfare implications of the wage replacement program are affected only very little by changing these two parameters.

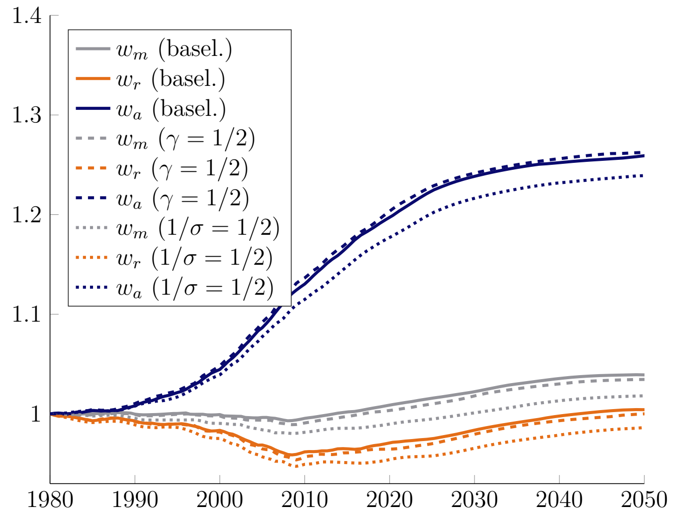

Figure 23: Wages per efficiency unit of labor

Notes: All lines show baseline transitions without policy. Solid lines: baseline calibration. Dashed lines: lower elasticity of labor supply. Dotted lines: lower elasticity of intertemporal substitution. All wages in occupation j are relative to w_j in 1980 in the respective calibration.

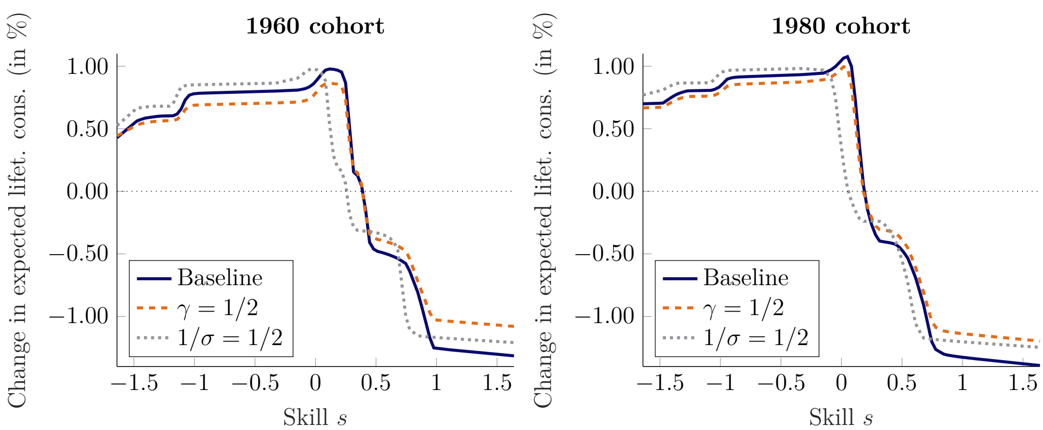

Figure 24: Expected lifetime consumption-equivalent welfare change of newborn households wage replacement of 420$

Notes: $(\bar{\phi}_{\tilde{T}}(s) - 1) \cdot 100\%$ is on the y-axis, where $\bar{\phi}_{\tilde{T}}(s)$ is as defined in equation (6).

G Policies Targeted at Switchers to Manual Occupation

In Section 6 the policies were targeted only at experienced routine workers who switched to the abstract occupation. In this section, I assume instead that experienced workers switching to the manual occupation become eligible to the wage replacement (loan).

Figure 25 plots average welfare gains from introducing the policies. Evidently, implementing the programs for switchers to manual does not increase welfare. Decomposing welfare changes as in the main text, for $L = 10,000\$$ and $R = 420\$$, reveals why this is the case. While at least in the case of the wage replacement program the partial equilibrium effect on welfare is positive (graphs not shown), the general equilibrium effects of both programs on welfare are negative.

There are several reasons for small welfare and efficiency gains when incentivizing experienced routine workers to switch to the manual occupation. First, the manual wage has risen only marginally relative to the routine wage since the 1980s, especially when one contrasts this to the rise in the abstract to routine wage. It is therefore rarely the case that foregoing their occupation-specific human capital in the routine occupation by switching to the manual occupation is optimal for experienced routine workers. Second, the calibration of the occupation-specific slopes a_j imply that many medium-skilled workers bear much lower wage losses when switching to the manual than when switching to the abstract occupation ($a_r - a_m < a_a - a_r$; as in Figure 2). The elasticity of manual labor supply with respect to the routine wage is therefore much higher than that of abstract labor supply.

For both of these reasons, the increase in the manual employment share, which is quantitatively small in any case, mostly runs via newborn or inexperienced households moving into manual occupations and less via experienced routine workers switching there in the baseline transition. Hence, the introduction of the policies mostly causes inefficient switching of the kind discussed in Section 6.2.1. Indeed, output does not increase on average under policies (graph not shown), in contrast to the case where policies where targeted at switchers to the abstract occupation.

Lastly note, however, that the model understates somewhat the overall rise in the manual relative to the routine wage over time, compared to the data (see discussion in Appendix D). This biases my results against finding significant welfare effects for policies that target switchers to the manual occupation.

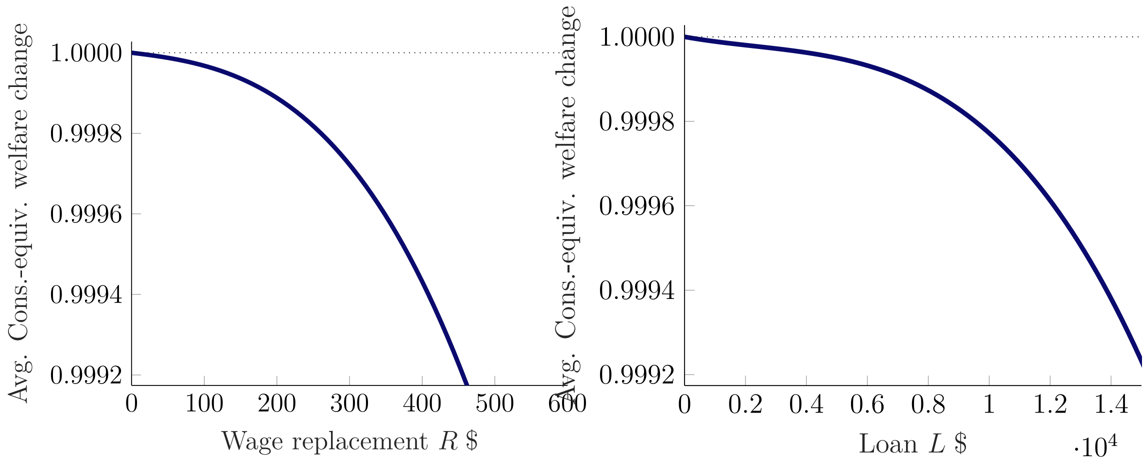

Figure 25: Average change in welfare under different policies (targeted at switchers from routine to manual) compared to the baseline.

Notes: Cons.-equiv. welfare change ϕ_t (defined in equation (5)), averaged between $t = 1980$ and $t = 2070$. Nominal quantities are expressed in 2017 US dollars.

H Additional Graphs and Tables

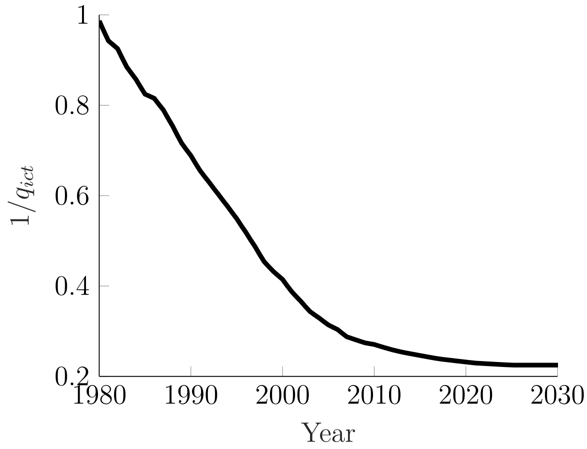

Figure 26: Relative price of ICT capital.

Notes: Data from 1980 to 2013 are taken from Eden and Gaggl (2018). I assume that $\frac{1}{q_{ict}}$ continues its fall at an average rate of 1% between 2013 and 2025, and stays constant thereafter.

Table 16: Parameters and targets for the calibration of Φ_ϵ.

Moment	Data	Model
Variance: ann. log earnings	0.70	0.70
Variance: 1-year change	0.26	0.26
Kurtosis: 1-year change	14.9	15.0
Fraction 1-year log change < 0.2	66.5%	70.2%
Fraction 1-year log change > 1.0	6.6%	6.6%

Notes: Empirical moments (Data) are taken from Guvenen et al. (2021), except for the variance of annual log earnings, which is not reported there. Instead, I use the value reported in Kaplan, Moll, et al. (2018). The moments simulated from the model only capture the exogenous component of earnings. Note that there is also an endogenous component due to the labor supply and occupational choice. I follow Kaplan, Moll, et al. (2018) in shrinking the calibrated grid of the productivity process ϵ by a factor of 1.85 to account for the endogenous labor supply choice. I also validate ex post that the earnings moments reported here are close to the ones implied by the full model.

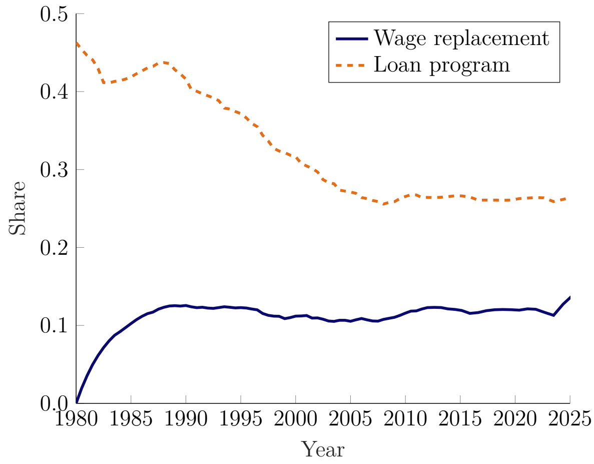

Figure 27: Share of eligible workers who work in the routine occupation under constrained optimal programs.

Notes: The figure shows the share of program participants who have returned to working in the routine occupation among all currently entitled households at time t. Averaged across the years 1980 to 2025 this share is 33% under the optimal loan program and only 11% under the wage replacement. This indicates a significant resource cost of the loan program, caused by workers who switch to the abstract occupation with no intent of working there in the long run.

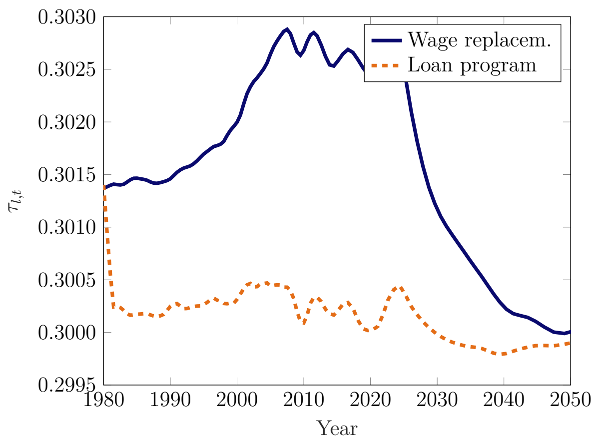

Figure 28: Labor income tax $\tau_{l,t}$ necessary to finance the policies.