

WORKER REALLOCATION ACROSS OCCUPATIONS: CONFRONTING DATA WITH THEORY

ETIENNE LALÉ

*University of Bristol
and IZA*

ABSTRACT. This paper studies the secular behavior of worker reallocation across occupations in the U.S. labor market. In the empirical analysis, we use microdata to construct consistent time series over a forty-five year period, and document that the fraction of employment reallocated annually across occupations is highly stable in the long run. We go beyond description and use an equilibrium model to identify potential changes in the productivity shocks and mobility costs that govern worker reallocation across occupations. We uncover the joint evolution of these factors by deriving a simple mapping between data and the model. Our analysis shows that constant reallocation rates across occupations mask slow-moving increases in the volatility of productivity shocks since the mid-1980s, and a pronounced upward shift in the cost of switching occupations in the period surrounding the Great Recession.

Keywords: Occupations, Reallocation, Wages, Equilibrium Search

JEL Codes: E24, J21, J31, J62

1. INTRODUCTION

There is an ongoing debate on whether labor markets have become more turbulent over the past half-century. This debate is to a large extent fueled by empirical studies that document time series of worker reallocation across, e.g., occupations or industries. Indeed, a standard view since the essays collected in Phelps et al. (1970) posits that the workforce is distributed over a range of distinct “islands” and reshuffles across them in response to island-specific productivity shocks. In this metaphor, more turbulent times should materialize through increased reallocation across segments of the labor market. The empirical evidence to date provides mixed results as there are trends in both directions, depending on the time series considered.¹ More importantly, the reason why the search for increased turbulence remains inconclusive is that the rates of worker reallocation may also be driven by other factors. Constant or even declining rates of reallocation could emerge in times of economic turbulence if there are counteracting changes in the other factors that affect these rates. One such example are changes in the cost of switching occupations, which are not easily controlled for because this cost is typically unobserved. Guidance from a model is, in this respect, paramount to interpret the patterns of worker reallocation found in the data.

Current version: October 2016. An online appendix is available at the web address: <http://www.efm.bris.ac.uk/e113851/papers/APPislands.pdf>. I am grateful to Iourii Manovskii and Etienne Wasmer for discussions that have greatly influenced this work. I also thank Christopher Flinn and two anonymous referees for their constructive comments that helped to improve the paper. All errors are my own.

Address: Department of Economics, University of Bristol, Priory Road, Bristol BS8 1TU, United Kingdom – Phone: +44(0)117 331 7912 – E-mail: etienne.lale@bristol.ac.uk.

¹For instance, on the one hand Kambourov and Manovskii (2008) document an increase in worker mobility across occupations and industries over the years 1968-1997. Davis, Faberman and Haltiwanger (2006), on the other hand, report a fall in job destruction rates and in the gross flows between unemployment and employment since the 1980s. Davis (2008) shows that the risk of unwanted job losses declined sharply in the U.S. during the same period.

In this paper, we contribute to this line of research at two levels. First, we construct new time series to analyze worker reallocation across occupations in the U.S. labor market. Relative to existing work, our time series exhibit several strengths, which we detail momentarily, and they cover a recent period that includes the Great Recession. Second, we map these data onto the parameters of an equilibrium model of worker reallocation that embodies productivity shocks and mobility costs. The mapping we establish allows us to disentangle the role of these two components in explaining the empirical patterns shown in the paper. We now provide details on our contributions.

The empirical sections of the paper document the evolution of net reallocation and excess reallocation from 1970 till 2015. Net reallocation measures the reshuffling required to accommodate changes in employment across occupations between two consecutive periods, ignoring the moves that cancel out in the aggregate.² Excess reallocation measures the latter, i.e. it is the difference between the total number of occupational switches and net reallocation.³ To our best knowledge, the behavior of these allocation processes in the U.S. has been documented in only two papers: Moscarini and Vella (2003), using data from the March Current Population Survey (henceforth March CPS) and Kambourov and Manovskii (2008), who use the Panel Study of Income Dynamics (henceforth PSID). Our estimates of net reallocation benefit from the much larger sample size of the March CPS relative to the PSID: every March CPS file provides us with a cross section representative of the U.S. population that allows to compute the employment share of each occupation even at a high level of disaggregation. In constructing our time series of excess reallocation, we take account of a number of pitfalls of the March CPS which have been pointed out by Kambourov and Manovskii (2013), and that imparted a substantial bias in previous estimates based on these data.⁴

In the subsequent step of the analysis, we use an equilibrium model to analyze the patterns of reallocation across occupations found in the data. We resort to the framework of Lucas and Prescott (1974) as it offers a direct formalization of the island parable put forward in the opening paragraph (Gallipoli and Pelloni, 2014). Recent research that has revamped this model, moreover, finds that it provides a basis for sound quantitative predictions; see, e.g., Alvarez and Shimer (2011). As in the study by Coen-Pirani (2010), we consider a version of the island model with gross flows that can be disentangled from net worker flows. We establish that, in the context of that model, occupational wages allow to estimate the parameters of the stochastic productivity process that drives worker flows in excess of net reallocation flows. Accordingly, we estimate these parameters using wage data, we feed them into the model alongside our estimates of net reallocation, and then we recover mobility

²Formally, net reallocation is defined as the sum of the absolute changes in occupational employment shares, divided by two to adjust for double counting (Section 2). Murphy and Topel (1987) and Layard et al. (2005) use this measurement (applied to industry employment shares) in relation to the study of unemployment.

³Excess reallocation is often referred to as “churning”; see Moscarini (2001). Jovanovic and Moffitt (1990) emphasize the importance of studying net reallocation and excess reallocation jointly. They interpret net reallocation as stemming from shifting demands across different segments of the labor market, which is emphasized by Lucas and Prescott (1974), Lilien (1982) and recently by Kambourov and Manovskii (2009) and Alvarez and Shimer (2011). In contrast, excess reallocation is supposed to result from idiosyncratic uncertainty at the job-match level rather than economy-wide changes. According to Jovanovic and Moffitt (1990), this in turn would explain why most of these moves cancel out.

⁴We find excess reallocation rates that are more than twice higher than those tabulated by Moscarini and Vella (2003). Appendix B discusses the underlying measurement issues at length. Another difference with Moscarini and Vella (2003) is that we use occupational classifications the categories of which are invariant over the entire period examined. So doing, we avoid several breaks and inconsistencies in the time series derived from these classifications.

costs which are pinned down by excess reallocation in the equilibrium of the model. By applying this procedure to each decade of the period under study, we quantify the importance of productivity shocks and mobility costs in explaining worker reallocation subperiod by subperiod. To be precise, we are able to identify and interpret *changes* in the levels of mobility costs across periods. We note that, at the same time, there are some limitations in using the model to interpret the *levels* of mobility costs. In sum, this semi-structural approach allows us to draw inferences on whether the role played by productivity shocks and mobility costs has evolved over time.

The findings of the paper are as follows. First, we document that worker reallocation across occupations has been remarkably stable since 1970. Over the period considered, the rates of net reallocation across (3-digit) occupations have remained around 4.4 percent and those of excess reallocation at 14.6 percent per year. In line with Kambourov and Manovskii (2008), we find a mild increase in net reallocation between the mid-1980s and mid-1990s. We also find that this was reverted in the 10-year period that followed. Second, generally there has been a slow-moving increase in the volatility of productivity shocks during the period 1976–2015, albeit with an interruption between 1996 and 2005. Excess reallocation has been slightly higher in more recent decades too, but viewed through the lens of the model, mobility costs have remained rather steady in the long run. Third and conversely, the last decade stands out by displaying much higher volatility of productivity shocks, which supports the idea of *increased* turbulence. Meanwhile, the rates of excess reallocation in the last decade were not too different from those previously observed. The model therefore implies that the increase in turbulence was accompanied by an increase in the costs of moving to a different occupation. These changes may have been felt at the level of occupation-industry cells: indeed, we obtain a similar picture when the model is employed to study worker reallocation across industries.

This paper is related to a strand of literature concerned with economic turbulence and its implications for labor markets. The term “economic turbulence” was coined by Ljungqvist and Sargent (1998) to denote the idea that changes in the macro-environment (e.g., the advent of new technologies) may result in more disruptive labor market trajectories. This view is not undisputed, however, and there are also controversies as to the labor market implications of increased turbulence. For instance, Ljungqvist and Sargent (1998, 2008) argue that the high European unemployment rates are a consequence of more turbulent times, whereas in den Haan, Haefke and Ramey (2005) turbulence leads to lower unemployment. Other examples include Fujita (2015) who studies the secular decline of the separation rate in the U.S. labor market, and Lalé (2016) who studies employment trends among older workers on the two sides of the Atlantic. The approach we take is different from that in these papers. We study worker reallocation through the lens of a model which has unambiguous predictions as to the labor market implications of turbulence. We do not posit that turbulence has increased; instead, we use a wage equation derived from the model to confront this hypothesis.

The analysis also contributes to research that uses the island model of Lucas and Prescott (1974) as a tool for quantitative investigations. Prominent examples include Alvarez and Veracierto (2000, 2012) and Kambourov and Manovskii (2008) who analyze, respectively, the effects of labor market policies and the link between human capital and wage inequality. In the migration literature, the island model has been used to study the behavior of worker flows across U.S. states (Coen-Pirani,

2010), the dispersion of house prices across metropolitan areas (Van Nieuwerburgh and Weill, 2010) or the dynamics of migration in and out of cities following a productivity shock (Davis, Fisher and Veracierto, 2016). As already mentioned, we relate our work to Coen-Pirani (2010) who demonstrates how the joint behavior of gross and net flows is informative as to the underlying allocation process. There is also a relationship between this paper and the study by Alvarez and Shimer (2011). They develop a continuous-time, tractable version of the island model, wherein one can obtain a mapping between industry-level wages and unemployment. This is similar in spirit to the mapping we derive between wages and the process for productivity shocks in the model. The difference is that they seek to uncover the parameters of a regulated Brownian motion (recall that their model is set in continuous time) whereas our mapping is for a discrete-time, mean-reverting process. Further, their focus is on unemployment whereas we interpret reallocation as net mobility across occupations and we do not study whether this is mediated by a spell of unemployment.^{5 6}

Finally, although we do not study business-cycle fluctuations explicitly, there seems to be a link between our findings for the period that includes the Great Recession and the literature on mismatch unemployment along the lines of Şahin et al. (2014). The authors build a measurement framework that bears resemblances with the island metaphor of the labor market. They find that mismatch across industries and occupation plays a limited role in explaining the increase in U.S. unemployment in the recession of 2007-09, yet that it has increased during this period. Barnichon and Figura (2015) and Herz and Van Rens (2015) corroborate these results. Also, Herz and Van Rens estimate that worker mobility costs have remained almost constant during the Great Recession, which differs from our finding for the most recent period. We view our result as complementary to theirs because they use a much different measurement framework and different method to infer mobility costs.⁷ Our estimate of an increase in mobility costs for the period starting in 2006 could in turn explain higher mismatch across occupations during the Great Recession.

The paper is organized as follows. Section 2 documents the long-run evolution of worker reallocation across occupations in the U.S. Section 3 presents an outline of the model and studies the relationships between reallocation, productivity shocks and mobility costs. In Section 4, we explain our methodology to connect data to the model, which we apply in Section 5. Section 6 concludes.

2. DATA: WORKER REALLOCATION ACROSS OCCUPATIONS SINCE 1970

This section presents the data and measurement of worker reallocation across occupations, and then documents its evolution for the period from 1970 onwards.

⁵The relationship between unemployment and reallocation across industries as in the island model of Lucas and Prescott (1974) is at the heart of the sectoral shift hypothesis studied by Lilien (1982), and discussed in a subsequent paper by Abraham and Katz (1986). See Gallipoli and Pelloni (2014) for an overview of this debate.

⁶The theme of unemployment is also pursued by Carrillo-Tudela and Visschers (2013), Lkhagvasuren (2012) and Wiczner (2015); they develop computationally tractable variants of the island model.

⁷To be precise, Herz and Van Rens (2015) consider an environment with a continuum of submarkets, each of which has a Diamond-Mortensen-Pissarides structure. In every submarket, workers and employers bargain on wages to split the surplus from matching. Thus, the authors allow for multiple sources of mismatch, such as costs of worker mobility and frictions in wage determination. On the other hand, in the model that we consider, wages are set competitively and there is no friction in the matching process within the island. Therefore the mobility costs we estimate embed a form of labor market friction that is distinct from that in their paper.

2.1. Data and Measurement⁸. Our primary measurement of worker reallocation is *net* reallocation across occupations. Letting $\pi_{o,t}$ denote the share of aggregate employment in occupation o at time t , net reallocation over a one-year period is defined as:

$$(1) \quad \text{net}_t = \frac{1}{2} \sum_o |\pi_{o,t} - \pi_{o,t-1}|.$$

This statistics measures the reshuffling of employment across occupations between year $t - 1$ and t , net of those worker flows that cancel out.⁹ The definition is standard and is often analyzed in conjunction with gross worker flows:

$$(2) \quad \text{gross}_t = \frac{1}{N_t} \sum_i \mathbb{1}\{o_{i,t} \neq o_{i,t-1}\}.$$

In equation (2), $o_{i,t}$ is the occupation of employment of worker i in year t , N_t is the number of workers employed both at $t - 1$ and t and $\mathbb{1}\{\cdot\}$ is the indicator function. *Excess* reallocation is defined as the difference $\text{gross}_t - \text{net}_t$, which measures the flows that cancel out in the aggregate.¹⁰

Data Sources. In equation (1), the measurement of net reallocation relies on estimates of the occupational employment shares $\pi_{o,t}$ in each year t . Therefore, it is clear that it requires “only” a series of cross sections representative of the population under study. We use the Integrated Public Use Microdata Series (IPUMS) (<http://cps.ipums.org/cps/>) collection of the March CPS to construct the time series of net reallocation. The reason why the March CPS is well suited for this purpose is twofold. First, the large size of each cross section helps avoid small-cell problems when looking at reallocation at a high level of disaggregation. Second, in the March CPS, a respondent reports detailed information about her employment status in the year prior to the survey. The occupational (and industry) affiliation which is part of this information is generally considered less noisy compared to, say, occupations in the monthly files of the CPS.

There are additional appealing features of the CPS files made available by the IPUMS-CPS project. Indeed, to enhance the comparability of occupational data in historical U.S. Census samples, the IPUMS-CPS project recoded occupations according to different classification schemes that remain consistent over the years of the period considered. We rely throughout on the so-called OCC1990 classification developed by Meyer and Osborne (2005) at the Bureau of Labor Statistics. This classification contains 7 categories at the 1-digit level, 80 at the 2-digit level and 387 categories at the 3-digit level. The online appendix provides further details about these data.

Next, to construct gross worker flows, we combine two data sources. The first is the monthly files of the CPS Outgoing Rotation Group samples (available at: <http://www.nber.org/morg/annual/>). In these files, we link respondents longitudinally to construct gross flows over a one-year period. One drawback of this approach is that these flows are likely contaminated by measurement error. To overcome this issue, we use the Job Tenure and Occupational Mobility supplements of

⁸We refer the reader to Appendix B of the paper for a detailed discussion of the relevant measurement issues, and to the online appendix for additional information on the data and the sample used in the analysis.

⁹Notice that the occupational employment shares $\pi_{o,t}$ (resp. $\pi_{o,t-1}$) are defined with respect to aggregate employment in period t (resp. $t - 1$). For all t , $\sum_o \pi_{o,t} = 1$.

¹⁰Thus, the statistics gross_t and net_t would coincide if, for every occupation-cell in year t , no worker joins an occupation when at least one worker leaves that occupation and vice versa.

the CPS (<http://www.nber.org/data/current-population-survey-data.html>). These are administered biennially, and thus they provide us only with discrete snapshots. Meanwhile, in these supplements occupations in year $t - 1$ and year t are recorded using dependent coding, which is a more reliable method of data collection (see Appendix B). We adjust the gross flows based on the Outgoing Rotation Group samples to match the figures derived from the Occupational Mobility supplements in the overlapping periods. Prior to making these adjustments, we recode occupations using the time-invariant classification from the IPUMS-CPS.

Measurement Issues. We remark briefly on some measurement issues here. First, we must define aggregate employment. We circumscribe it to the population of civilians of working age who are not self-employed, employed in a family business or working for the government. In the online appendix, we show that these sample restrictions are inconsequential for our results. Second, in equations (1) and (2), we must decide whether employment refers to the number of employed persons or to the total number of hours they work. We report results for both. We prefer the first measurement because it relates directly to the model in the next section, which does not have an intensive margin.

The other remark concerns some shortcomings of the March CPS pointed out by Kambourov and Manovskii (2013). In Appendix B, we summarize their arguments and discuss whether (and how) these could affect our results. In short, the measurement of *net* reallocation does not rely on individual transitions across occupations, and thereby it is immune to most of the shortcomings identified by the authors. For *excess* reallocation, we corroborate the findings of Kambourov and Manovskii (2013), that the March CPS cannot be used to measure gross worker flows over a one-year period. As just discussed, we rely on other data sources to circumvent this problem.

2.2. Baseline Figures.

Net Reallocation. The charts in Figure 1 display net reallocation at the different digit levels of the occupational classification. The fraction of total employment which is reallocated across 1-digit occupations between two consecutive years is slightly less than 1 percent. That figure increases to 2.4 percent at the 2-digit level and to 4.4 percent at the 3-digit level. The figures are virtually identical when employment is weighted by the number of hours worked by employed individuals. Worker reallocation at the 1- and 2-digit levels is remarkably stable over the entire period considered. The 3-digit level, on the other hand, exhibits an upward trend in the 1980s and early 1990s, which is reverted in the period beginning in 1995. Finally, there is no apparent relationship between the different time series and the recessionary periods covered by the data.

To put these findings in perspective, we compare them with the figures reported by Kambourov and Manovskii (2008). The first difference relates to the levels of net reallocation reported in their paper and in Figure 1: they find significantly higher levels at all digit levels. We show in the online appendix that a large part of the discrepancy between their estimates and ours comes from differences in sample dispositions. Other factors that may play a role include differences in the occupational classification and the smaller size of the PSID sample. We note that, on the other hand, Moscarini and Vella (2003) report an average of 2 percent for net reallocation at the 3-digit level. This figure is far off from our results and from estimates based on the PSID.

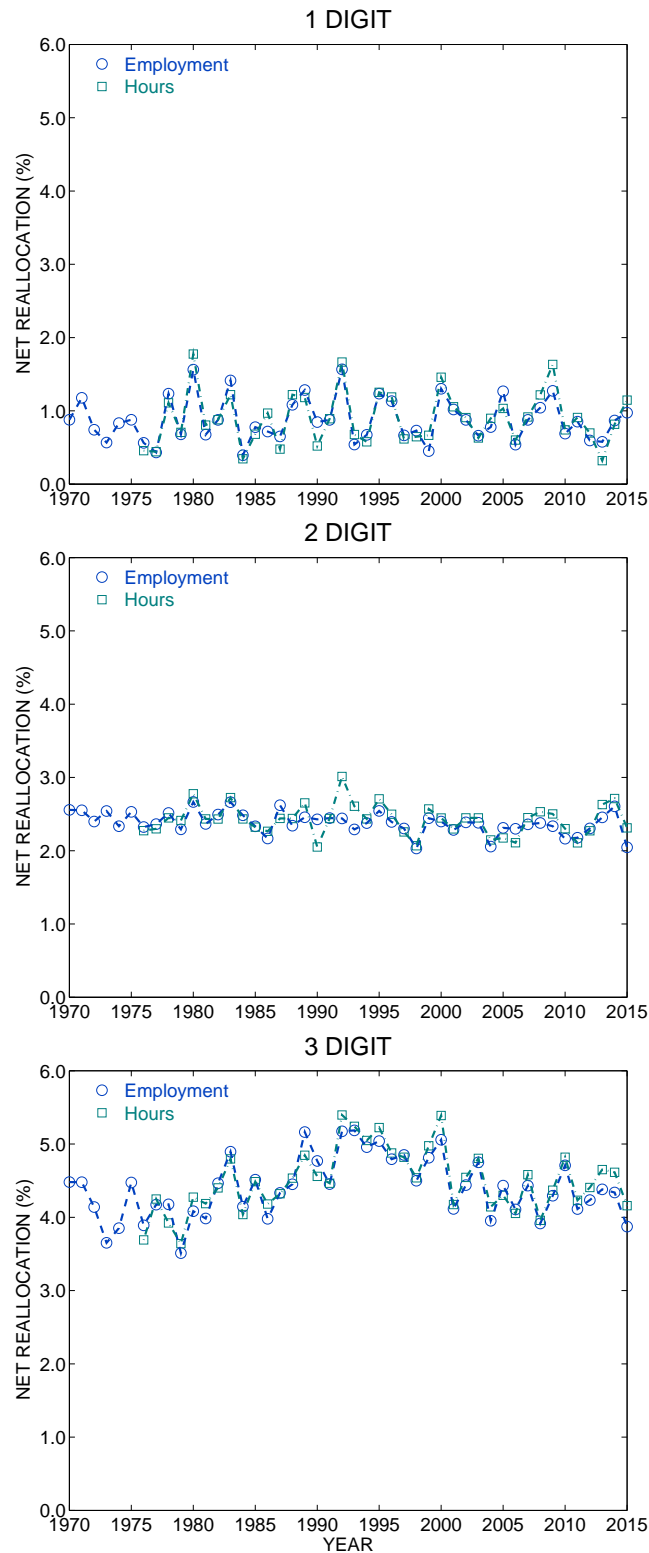


Figure 1. Net reallocation across occupations: Baseline figures
 The upper, middle and lower charts display, respectively, net reallocation rates at the 1-, 2- and 3-digit level of the occupational classification. Circles and squares denote, respectively, employment-weighted and hours-weighted rates of net reallocation. The hours variable is not available prior to 1976, and so the hours-weighted time series are only from 1976 onwards.

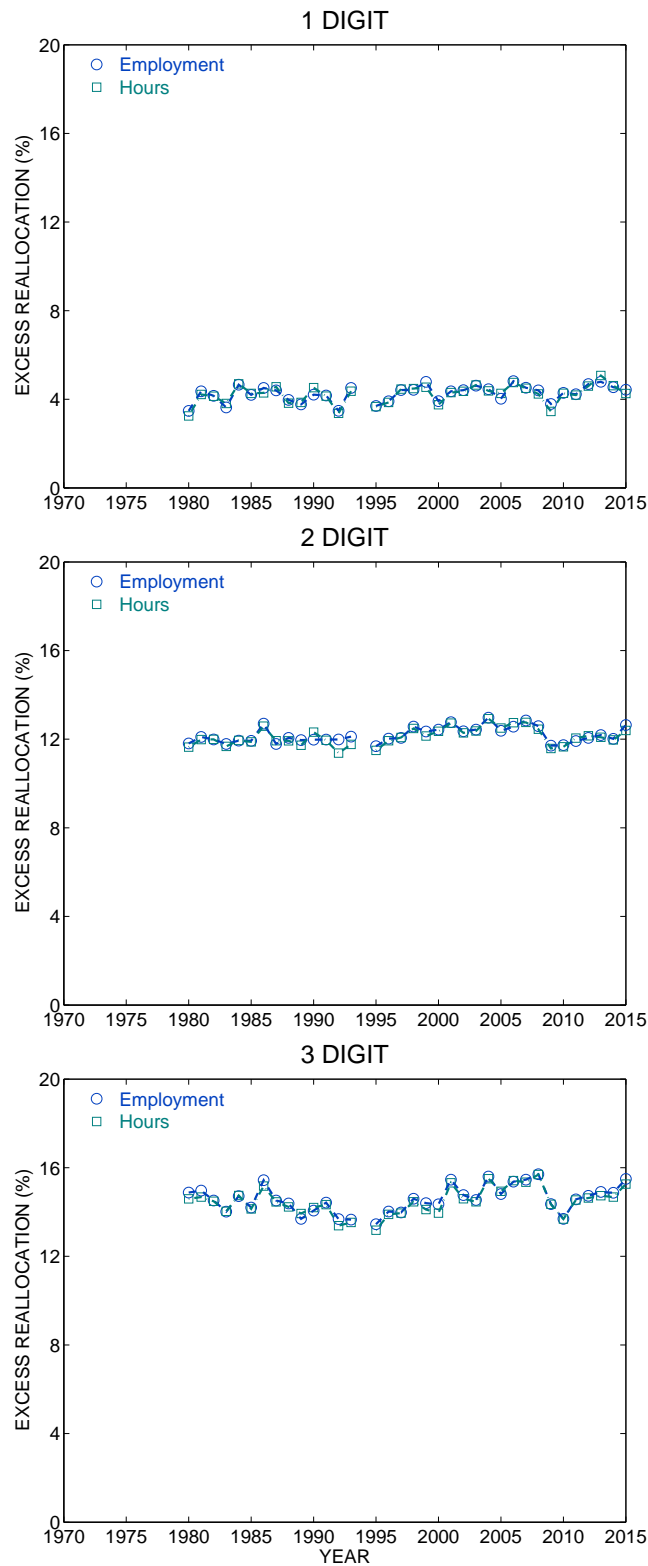


Figure 2. Excess reallocation across occupations: Baseline figures
 The upper, middle and lower charts display, respectively, excess reallocation rates at the 1-, 2- and 3-digit level of the occupational classification. Circles and squares denote, respectively, employment-weighted and hours-weighted rates of excess reallocation. The time series begin in 1980 which is the first period of observation for gross occupational mobility (see Appendix B.2).

We now comment on the trends displayed in Figure 1. Kambourov and Manovskii (2008) report positive trends in net reallocation for the years 1970-1997, notably at the 3-digit level of the occupational classification. The 3-digit level in Figure 1 confirms their findings for the overlapping period: we observe a rise in net reallocation until the mid-1990s. When looking at the data decade by decade, we find, moreover, that the increase is statistically significant, and that it was matched by a decrease of similar magnitude in the subsequent period (see Table 1 below). In Section 5 of the paper, these changes contribute to inform the model used to analyze excess worker reallocation.

Excess Reallocation. Figure 2 shows the evolution of excess reallocation across occupations since 1980. We find that excess worker reallocation is on average 4.3 percent at the 1-digit level, 12.2 percent at the 2-digit level, and finally 14.6 percent at the 3-digit level. In other words, out of all 3-digit occupational moves that occur over a one-year period, we find that about three quarters (14.6 percent divided by gross occupational mobility, which averages at 19.0 percent) of these moves cancel out. As previously mentioned, this proportion is typically interpreted as measuring the role of churning in generating reallocation across occupations.

The figures we report for the importance of excess reallocation (measured by the ratio with respect to overall reallocation) are in the upper range of available estimates. In the PSID data analyzed by Kambourov and Manovskii (2008), for instance, excess reallocation accounts for two thirds of reallocation across occupations. In the online appendix, we show that this difference is to a large extent driven by the fact that our sample includes younger workers, who are more mobile across occupations. Interestingly, we find that churning plays a quantitatively similar role in explaining reallocation across 3-digit *industries* (see Subsection 5.3 and Appendix A.2): our estimates indicate a contribution of excess reallocation by 70 percent.

Turning to changes across periods, a close look at Figure 2 suggests that: (i) excess reallocation has been stable in the 1980s and early 1990s, (ii) that it has been more volatile and on an upward course during the late 1990s and 2000s, and (iii) that it has dropped during the period surrounding the Great Recession. It is worth noticing that the slight increase of excess reallocation coincides with the decrease in net reallocation between 1995 and 2000. That trend came to a halt during the 2000s. This fact is consistent with the decrease in occupational mobility documented by Moscarini and Thomsson (2007) and Moscarini and Vella (2008), and more generally with the downward trend in job-to-job mobility in the U.S. labor market (see, e.g., Hyatt and Spletzer, 2013).

2.3. Robustness. The online appendix provides a host of robustness checks. First, we analyze net reallocation and excess reallocation computed using various sample restrictions. In one instance, we show that the trends are similar when employment is restricted to male workers. Thus, although major changes affected women's labor participation during the period considered, it does not seem that the trends in net reallocation are driven by new female workers who would predominantly direct themselves to particular occupations. In another instance, we show that the results with respect to net reallocation are robust to restricting the sample to prime-age workers. This suggests that entries to and exits from the labor market are not concentrated into specific occupations, and therefore that they do not explain the trends displayed in Figure 1. As already mentioned, this sample restriction

decreases the levels of excess reallocation substantially compared to Figure 2. Meanwhile, it has no impact on the changes across periods discussed in the previous subsection.

Second, we show that, under similar sample dispositions, our estimates of net and excess reallocation are consistent with those of Kambourov and Manovskii (2008) based on PSID data. Thus, our results are not driven by the limitations of the CPS uncovered by Kambourov and Manovskii (2013).

Third, to gain a better understanding of the process captured by our measure of net reallocation, we recompute this time series using a longer time lag. With a time horizon of 5 years (vs. 1 year in equation (1)), we find that the rates of net reallocation are of course higher, yet that they do not increase in proportion to the length of the time window. We will argue in the next section that this lines up well with reallocation in the model of Lucas and Prescott (1974), wherein islands are randomly losing or gaining workers every period. In addition, we show that the long-run stability of net reallocation rates is robust to lengthening the time window to define these rates. These findings hold true at the different digit levels of the occupational classification.

2.4. Towards the Model. In order to draw inferences, in the remainder of the paper we employ Lucas and Prescott (1974)'s island model as a quantitative tool to analyze the patterns just described. We think that there is room to discuss whether this is the “right” framework to decode the data. Accordingly, before introducing the model, we examine the arguments in support of this choice. Table 1 summarizes the empirical patterns that the model will help understand.

Table 1. Worker reallocation across occupations: Levels and trends

	1976-2015	By subperiod			
		1976-1985	1986-1995	1996-2005	2006-2015
A. Net reallocation					
Level	4.44 (0.07)	4.18 (0.12)	4.75 (0.13)	4.57 (0.11)	4.24 (0.08)
Trend	0.15 (0.58)	7.13 (3.62)	10.48 (3.27)	-6.26 (3.40)	-0.72 (2.91)
B. Excess reallocation					
Level	14.6 (0.10)	14.6 (0.15)	14.1 (0.19)	14.7 (0.18)	14.9 (0.22)
Trend	1.83 (0.95)	-13.3 (7.66)	-17.3 (4.63)	12.4 (4.53)	-3.07 (7.13)

NOTES: Net reallocation and excess reallocation across 3-digit occupations measured using the employment weights (circles, lower chart in Figures 1 and 2). The trend is computed using a linear regression of the time series against calendar years and the coefficient is multiplied by 100 for legibility. Standard errors in parentheses.

First, a robust feature of the data is that, between two consecutive years, some occupations experience employment gains, other occupations suffer some employment losses, and gross employment changes dwarf net employment changes. The crux of the island model is to capture the phenomenon of reallocation between expanding and declining islands, which makes it natural to choose this framework. To be precise, in the model analyzed by Lucas and Prescott (1974), one cannot distinguish

between gross flows and net flows. The model introduced in Section 3 is an amended version that allows for additional gross flows on top of the net worker flows. Second, the patterns we document are for a set of invariant occupations and, later on in the analysis, we focus on occupations that display positive levels of employment in each year of the period considered. This lines up well with another feature of the model, that no island enters or exits the market. Further, the Inada condition in the model prevents islands from shrinking to a zero level of employment. Third, the model embodies two simple economic forces that drive worker reallocation: productivity shocks and mobility costs. As argued in the introduction, both are the object of recent scrutiny in the literature. Fourth, this is an equilibrium model with very few parameters. It is also parsimonious in that it abstracts from fluctuations at the intensive margin. As seen in Figures 1 and 2, the patterns of reallocation across occupations are virtually identical when individual weights are adjusted to account for hours worked.

3. THEORY: AN EQUILIBRIUM SEARCH MODEL

This section introduces the theory which we use to confront the data. In the first subsection, we describe the key features of Lucas and Prescott (1974)'s model briefly and refer the reader to their paper for a detailed exposition. In the second subsection, we highlight some key implications of that model by using numerical examples.

3.1. The Model. Time is discrete and runs forever. There is a continuum of workers and a large number of distinct submarkets (islands). Each island engages in the production of a homogeneous good, which is constrained by an island-specific productivity level and a resource constraint stemming from exogenous and endogenous worker reallocation. The state of an island is given by its idiosyncratic productivity, z , and the number of workers in the island, ℓ .

A model period consists of two subperiods. In each island, during the first subperiod, an exogenous shock separates a fraction δ of workers from the island. Letting $\tilde{\ell} (= (1 - \delta)\ell)$ denote the number of workers who remain on the island after the exogenous shock, mobility decisions take place during the second subperiod and these result in an endogenous employment level, $n(z, \tilde{\ell})$. Production also occurs during the second subperiod. The production technology is:

$$(3) \quad f(z, n) = e^z n^\alpha$$

with $0 < \alpha < 1$, i.e. there are decreasing returns to scale. The wage w is determined competitively in each island, so that two constraints are satisfied in equilibrium:

$$(4) \quad \begin{aligned} w(z, \tilde{\ell}) &= \alpha e^z n(z, \tilde{\ell})^{\alpha-1} \\ n(z, \tilde{\ell}) &\leq \tilde{\ell}. \end{aligned}$$

The productivity level z fluctuates over time in response to shocks that are independent across islands. These shocks occur at the beginning of the period, and the law of motion of z is a Markov process with transition function $F(\cdot|z)$, i.e. $F(z'|z) = \Pr\{z_{t+1} < z' | z_t = z\}$. $F(\cdot|z)$ is assumed to be a decreasing function of z to generate persistence in productivity. We specialize $F(\cdot|z)$ later on in the analysis. According to the realizations of this process, workers who are still on the island after the δ shock decide to stay or to walk away. Workers are risk-neutral, and therefore this decision

is consistent with maximization of the discounted present value of their income streams. Workers discount the future with a factor $\beta \in (0, 1)$.

The mechanisms of reallocation rests on assumptions on how workers move across islands. Specifically, we assume that a worker who leaves an island (either exogenously or endogenously) incurs a mobility cost c , and is assigned to the island of her choosing at the beginning of the next period. We allow workers to be perfectly informed about the current state of every island, so that in equilibrium they direct their search towards expanding islands. We motivate this assumption with the parametrization of the model: a period is interpreted as one year, which in turn suggests that the search process should not last longer than one period. On the other hand, under random search, workers would sometimes land on the “wrong” (low productivity) island and would have to keep searching. Finally, notice that since a worker belongs to no island while searching, she receives no income during this period. Thus, the mobility cost c is paid in addition to the opportunity cost of search.¹¹

To formulate the problem of a worker on an island with current state (z, ℓ) , denote by $v(z, \ell)$ her value function at the beginning of the first subperiod, by $\tilde{v}(z, \ell)$ her value function at the beginning of the second subperiod, and by v_s her value at the beginning of the next period if she leaves the island.¹² Since a worker cannot move during the first subperiod, we have:

$$(5) \quad v(z, \ell) = \delta (\beta v_s - c) + (1 - \delta) \tilde{v}(z, (1 - \delta) \ell).$$

The value $\tilde{v}(z, \ell)$ satisfies: $\tilde{v}(z, \ell) = \max \{ \beta v_s - c, w(z, \ell) + \beta E(v(z', \ell')) \}$, where $E(\cdot)$ is the expectation with respect to z' and ℓ' conditional on z and ℓ .¹³

In an equilibrium, we must examine two possibilities.^{14 15} First, if the island is attractive (for example productivity is high), then some workers arrive on the island. Taking into account the attrition that occurs during the first subperiod, this equates $v(z, \ell)$ to $\delta (\beta v_s - c) + (1 - \delta) (w(z, (1 - \delta) \ell) + \beta v_s)$, i.e. the flow of new arrivals stops when workers expect that they will be indifferent between searching and being on the island at the beginning of the next period. The other possibility is that no worker arrives next period, which implies that the expected value $E(v(z', \ell))$ is not larger than v_s at the level ℓ that prevails during the second subperiod. Piecing the cases together yields:

$$(6) \quad \tilde{v}(z, \ell) = \max \left\{ \beta v_s - c, w(z, \ell) + \beta \min \left\{ v_s, \int v(z', \ell) dF(z'|z) \right\} \right\}.$$

Lucas and Prescott (1974) show that, for a given value of search v_s , equation (6) defines a contraction mapping for \tilde{v} when $\tilde{v} = v$, i.e. when $\delta = 0$. It can be checked that for $\delta > 0$ the system of (5)–(6) is a contraction mapping as well. Next, the value v_s is endogenous from the aggregate point of view because the average size of the labor force per island must be a fixed quantity. Intuitively, a high value

¹¹To anticipate on the numerical exercise, the cost of moving across islands, c , pins down the rate of excess reallocation in the equilibrium of the model.

¹² v_s denotes the value at the beginning of the next period of leaving either exogenously (i.e. during the first subperiod) or endogenously (during the second subperiod). These two values coincide because a worker who belongs to no island must wait until the beginning of the following period to be allocated to an island.

¹³In order to simplify notations, we write the asset value $\tilde{v}(z, \ell)$ as a function of ℓ instead of $\tilde{\ell}$. Obviously, in $\tilde{v}(z, \ell) = \max \{ \beta v_s - c, w(z, \ell) + \beta E(v(z', \ell')) \}$, the number of workers in the island at the start of the first subperiod was $\ell/(1-\delta)$.

¹⁴We consider steady-state equilibria, so that the two possibilities are for a stationary environment.

¹⁵The third possibility which we need not discuss here is that workers leave the island. This occurs when $w(z, \ell) + \beta E(v(z', \ell')) \leq \beta v_s - c$, so this case is taken care of by the max operator in the Bellman equation that defines $\tilde{v}(z, \ell)$.

of search induces too many workers to be mobile across islands and hence too small a labor force per island. Lucas and Prescott (1974) demonstrate that, for any fixed size of the labor force, there is a unique equilibrium value v_s .¹⁶

3.2. Illustration of the Workings of the Model. To gain an understanding of how the model generates excess worker reallocation, in this subsection we examine two numerical examples. In the first, we vary the level of the mobility cost c and in the second example we change the volatility of productivity shocks. Both examples are based on the benchmark parametrization (see Section 5).¹⁷

Mobility Costs. In the first example, we study the effect of the mobility cost c on the employment level within an island. A high mobility cost reduces the value of being in any island for a given value of search v_s , and makes workers less willing to leave an island. To be precise, workers join until the expected value of being on an island next period is equated to the value of search. Therefore, an island is attractive if its *beginning-of-period* labor force ℓ satisfies $\ell \leq \underline{\ell}(z)$, with $\underline{\ell}(z)$ implicitly defined by:

$$(7) \quad \int v(z', (1 - \delta)\underline{\ell}(z)) dF(z'|z) = v_s$$

On the other hand, workers endogenously leave the island when they are too numerous, i.e. the beginning-of-period ℓ is such that $\ell \geq \bar{\ell}(z)$ where $\bar{\ell}(z)$ solves:

$$(8) \quad w(z, (1 - \delta)\bar{\ell}(z)) + \beta \min \left\{ v_s, \int v(z', (1 - \delta)\bar{\ell}(z)) dF(z'|z) \right\} = \beta v_s - c$$

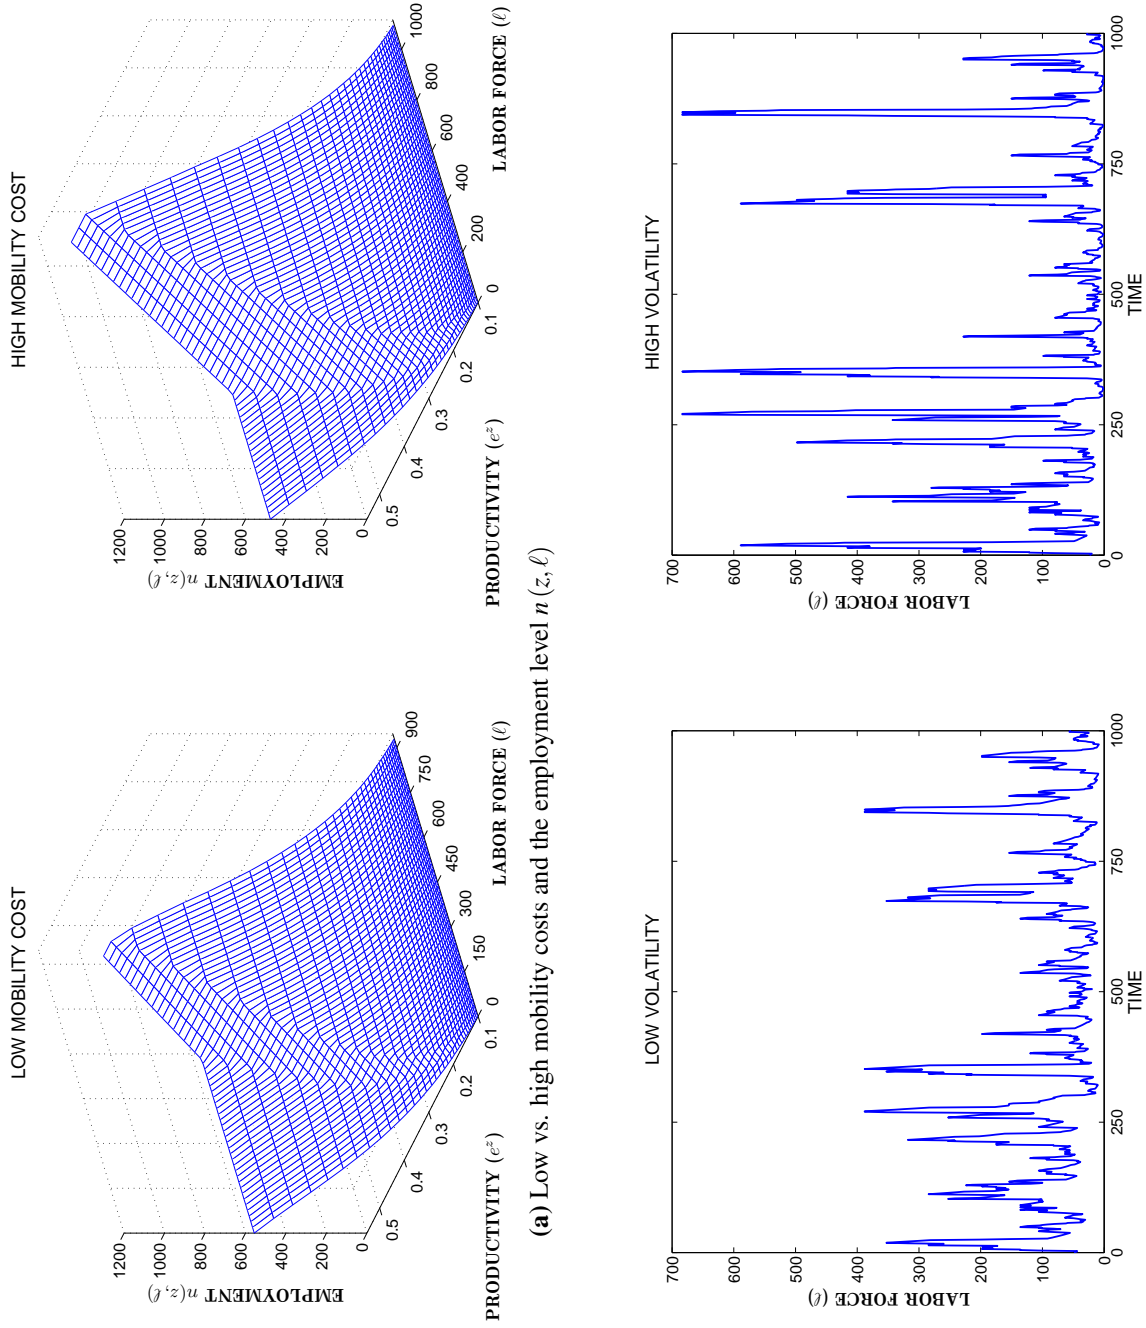
Finally, when $\underline{\ell}(z) \leq \ell \leq \bar{\ell}(z)$ the labor force in the island depreciates at rate δ . An implication of equations (7) and (8) is that a higher mobility cost c reduces $\underline{\ell}(z)$ (lower inflow) and increases $\bar{\ell}(z)$ (lower outflow), which increases the range of inaction with respect to the size of the labor force, ℓ .

The graphs in Figure 3a illustrate these effects. In an equilibrium with high mobility costs, the lower bound $\underline{\ell}(z)$ is reduced as shown in the low- ℓ high- z corner of the graphs (when the island is attractive). The fact that workers become less mobile, i.e. the increase in $\bar{\ell}(z)$, can be seen by looking at the upper kink along the ℓ dimension for each values of productivity z . Finally, observe that the ergodic set of labor force shown in the axis titled “labor force” spans a larger range of values when mobility costs are high.

Volatility of Productivity Shocks. The second example highlights the effect of more volatile shocks on the size of the labor force on an island. The key mechanism is that, in a more turbulent environment, productivity is more likely to jump from a low to high value and vice versa along the axis titled “productivity” in the graphs of Figure 3a. This in turn results in more frequent arrivals and departures in the island which raise the volatility of the size of its labor force.

¹⁶Notice that the equilibrium value of v_s incorporates the fact that after landing on an island, a worker may be returned to the search pool with probability δ . This feature is implicit in the comparison of $E(v(z', \ell))$ and v_s , which corresponds to the case where no worker arrives next period. As seen above, δ has also implications for the other case where some workers arrive on the island. The exogenous outflow increases the marginal productivity of those workers who remain on the island, which increases the wage and thus makes the island more attractive.

¹⁷The examples are constructed as follows. In the first one, we change the mobility cost by 50 percent below and above its baseline value. This results in excess reallocation rates of 18.5 percent and 11.5 percent, respectively (baseline excess reallocation is 14.5 percent). In the second scenario, we change the standard deviation of shocks to replicate these mobility rates. The low and high mobility cases are matched by setting σ to 0.142 and 0.238, respectively.



(a) Low vs. high mobility costs and the employment level $n(z, \ell)$

(b) Low vs. high volatility of shocks and the size of the labor force ℓ

Figure 3. Illustration of the workings of the model
 In the equilibrium with a low mobility cost (left, 3a) or with highly volatile shocks (right, 3b), excess reallocation is 18.7 percent. In the equilibrium with a high mobility cost (right, 3a) or with shocks of low volatility (left, 3b), excess reallocation is 11.5 percent.

To illustrate the mechanism, the graphs in Figure 3b display 1,000 period sample paths for the size of the labor force on an island.¹⁸ In a highly volatile economy, times of high productivity entail a larger increase in the productivity of labor. Therefore it takes a larger inflow of workers to offset the gap between the value of being on a productive island and the value of search. Hence the relatively higher spikes in the right graph of Figure 3b. Of course, in the data, yearly changes in employment within an occupation relative to its employment level are an order of magnitude lower than those depicted in the graphs. We think that Figure 3b is helpful in explaining the commonly-held view that an increase in turbulence should be associated with higher rates of reallocation across occupations.

4. CONNECTING THE DATA TO THE THEORY

Having presented the empirical patterns of interest and the model used to analyze them, we are in a position to explain how to connect the data to the theory.

First, we need to specify the stochastic process for productivity shocks. We adopt a standard specification, namely a first-order autoregressive process:

$$(9) \quad z' = \phi + \rho z + \sigma \varepsilon'$$

with: $\varepsilon \sim N(0, 1)$ and $0 < \rho < 1$. Henceforth, the number of parameters in the model is seven: the discount factor β , the curvature of the production function α , the separation probability δ , the mobility cost c , the mean ϕ , the persistence of the productivity process ρ and the standard deviation of innovations σ .

A simple relationship connects wages in the data to the parameters of the productivity process. The marginal product condition in (4) implies that wages in occupation o at time t satisfy:

$$(10) \quad \log(w)_{o,t} = \log(\alpha) + (\alpha - 1) \log(n)_{o,t} + z_{o,t}$$

after taking logs. Using equation (9) recursively, we have

$$(11) \quad \log(w)_{o,t} + (1 - \alpha) \log(n)_{o,t} = \phi + (1 - \rho) \log(\alpha) + \rho \left[\log(w)_{o,t-1} + (1 - \alpha) \log(n)_{o,t-1} \right] + \sigma \varepsilon_{o,t}.$$

The insight is that this last equation does not depend on the occupation-specific productivity level, z . Thus, after defining $y_{o,t}(\alpha) \equiv \log(w)_{o,t} + (1 - \alpha) \log(n)_{o,t}$, equation (11) suggests a family of auxiliary models to estimate the parameters ϕ , ρ , σ , namely:

$$(12) \quad y_{o,t}(\alpha) = \vartheta_0 + \vartheta_1 y_{o,t-1}(\alpha) + \vartheta_2 \varepsilon_{o,t}.$$

The parameters of interest are recovered as: $\widehat{\phi}_\alpha = \widehat{\vartheta}_0 - (1 - \widehat{\vartheta}_1) \log(\alpha)$, $\widehat{\rho}_\alpha = \widehat{\vartheta}_1$ and $\widehat{\sigma}_\alpha = \widehat{\vartheta}_2$. We maintain the subscript α on the parameters of the process to indicate that the estimates depend on the choice of α . We omit α in $\widehat{\vartheta}_0$, $\widehat{\vartheta}_1$, $\widehat{\vartheta}_2$ to save on notations, but naturally these parameters depend on the pre-specified value of α by construction of $y_{o,t}(\alpha)$.

¹⁸We use the same sequence of random numbers to feed the productivity process in the left and right graphs of Figure 3b – that is, the sequence of random integers that index the grid point for z . This improves legibility because this makes good shocks and bad shocks occur at the same time on the left and on the right graph. Since the volatility of shocks is different in the two cases, the *value* of z in good times and in bad times is of course different in the two scenarios.

Therefore the procedure to connect data to the model is as follows. First, we fix the discount factor using external information: we let $\beta = 0.951$ to accord with an annual interest rate of 5 percent. The model allows for one normalization, namely the average size of the labor force in each island.¹⁹ We set this number to 100. Second, we select a value for α and use data on wages and employment shares from the March CPS to construct the variable $y_{o,t}(\alpha)$. We arrange these data in a panel format and estimate the parameters ϕ , ρ , σ using the auxiliary model above. Specifically, we estimate equation (12) via maximum likelihood under the assumption of normality of the residuals.²⁰ Then we use the relationships: $\widehat{\phi}_\alpha = \widehat{\vartheta}_0 - (1 - \widehat{\vartheta}_1) \log(\alpha)$, $\widehat{\rho}_\alpha = \widehat{\vartheta}_1$, $\widehat{\sigma}_\alpha = \widehat{\vartheta}_2$. Finally, we plug these parameters into the model, set δ to the empirical value of net reallocation, and search the value of the cost c that aligns the model-generated level of excess reallocation to the data. In other words, the model matches *exactly* the empirical value of net reallocation, for exogenous reasons via δ , as well as the value of excess reallocation, endogenously via c . Appendix C provides details on our numerical methodology.

Parameters Uncovered by the Model. Before turning to the implementation, we highlight one caveat of the approach presented here. That is, we note that what equation (10) describes as the log-wage in an occupation does not have a well-defined empirical counterpart. There is hence some arbitrariness in the choice of this variable. In turn, the estimated parameters of the productivity process are also contingent on this choice. This suggests that excess reallocation and mobility costs as predicted by the model may not be directly interpretable. On the other hand, the model allows to interpret changes in mobility costs over time if we use the same definition of occupational wages across periods. This is the exercise we undertake in the next section.

5. APPLICATION: CHANGES IN THE DETERMINANTS OF WORKER REALLOCATION

We proceed in two steps to study the evolution of those factors that determine worker reallocation across occupations. First, we describe changes in the parameters of the productivity process which are indicated by the dynamic behavior of occupational wages. Then we feed the model with these parameters, use our estimates of net reallocation rates, and study their implications for the evolution of excess reallocation and mobility costs across periods.

One convention that we adopt throughout the analysis is to split the period 1976–2015 into four subperiods of equal length (10 years). Our main motivation comes from the contrast between the decades 1986–1995 and 1996–2005: as shown in Table 1, these are periods when the different series exhibit significant time trends and changes in the direction of the trend. This fact is also true for reallocation across industries (see Table A3 in the appendix). Finally, visual inspection of the time series of excess reallocation suggests that churning declined substantially during the Great Recession. Our model cannot speak to this feature of the data, but we can at least sense its effects by isolating the period surrounding this recession episode.

¹⁹Thus, $n_{o,t}$ in equations (10) and (11) can be replaced by the employment share $\pi_{o,t}$ defined in Section 2.

²⁰This is not a necessary assumption. In fact, the parameters ϕ , ρ , σ could also be recovered semi-parametrically by using the fact that the covariance between innovations in equation (9) must be equal to zero. The reason we maintain the assumption of normally-distributed residuals is only for consistency with the numerical solution of the model: in the computations we use Tauchen (1986)'s method to approximate the autoregressive process (9), and this method requires normality of the innovation term.

5.1. 1st Step: The Evolution of Productivity Shocks.

Methodology. To better align data with the model, we run a set of preliminary regressions to remove some layers of heterogeneity in wages. Specifically, we construct the residual of (log) hourly earnings for each individual by running the following OLS regressions:

$$(13) \quad \log(w)_{i,t} = x_{i,t}\zeta_t + v_{i,t}$$

for each year of the period under study.²¹ In this equation, $w_{i,t}$ denotes real hourly wages and $x_{i,t}$ includes a third-order polynomial in age interacted with educational dummies, marital dummies, race dummies, regional dummies and 1-digit occupational dummies.²² The residual log-wage of individual i in the time- t cross section is obtained using the predicted $\hat{\zeta}_t$ from the regression.

Empirical Results. Table 2 presents estimates of the parameters of the productivity process based on our benchmark specification: we use $\alpha = 0.675$ (a standard value for decreasing returns to scale in labor) and define occupational wages as the mean residual log-wage in the occupation.²³

Table 2. Estimates of the productivity process: Baseline specification

	1976-2015	By subperiod			
		1976-1985	1986-1995	1996-2005	2006-2015
$\hat{\phi}$	-0.107	-0.108	-0.103	-0.110	-0.110
	[-0.120,-0.095]	[-0.132,-0.083]	[-0.127,-0.078]	[-0.135,-0.084]	[-0.137,-0.083]
$\hat{\sigma}$	0.183	0.178	0.175	0.179	0.200
	[0.181,0.186]	[0.173,0.183]	[0.170,0.180]	[0.174,0.184]	[0.194,0.205]
$\hat{\rho}$	0.936	0.935	0.939	0.934	0.936
	[0.929,0.943]	[0.921,0.949]	[0.925,0.953]	[0.919,0.948]	[0.921,0.952]

NOTES: The table reports the estimates of the productivity process obtained using the mean (residual log-) wage in each occupation. The curvature parameter is $\alpha = 0.675$. Figures in brackets are 95 percent confidence intervals.

The productivity process exhibits high persistence in all periods, with the estimates of ρ always above 0.9. The estimates for the parameters ϕ and ρ appear stable across periods. For instance, the 95 percent confidence intervals overlap for most periods. On the other hand, we find a slightly U-shaped behavior for the standard deviation of shocks, σ : it decreases between the first two decades and then increases until the end of the sample period. In the next section, we use the model to assess the implications of this pattern.

²¹We experimented different ways of controlling for individual heterogeneity. An alternative to equation (13) is to pool the cross sections to run occupation-specific regressions, i.e. to estimate a vector ζ_o instead of ζ_t where o denotes 3-digit occupations. The results we obtained were similar to those presented in this section. Equation (13) is our preferred specification because it deals explicitly with changes in the returns to education during the period analyzed.

²²The interaction is between the polynomial in age and the educational dummies. The educational dummies are for “less than high school”, “some college” and “college or higher education”; the reference category is “high school graduates”. The marital dummies are for “separated, divorced, widowed” and “single or never married”; the reference category is “married”. The race dummies are for “Blacks” and “other”; the reference category is “White”. Regional dummies are for the nine standard regions and divisions of the United States.

²³In the estimation, we only use occupations with valid wage information over the whole 1976-2015 period; see the online appendix for details. So doing, we require the panel to be perfectly balanced.

5.2. 2nd Step: The Evolution of Mobility Costs. Table 3 presents the benchmark results: we report the value of excess reallocation that the model replicates and the mobility cost c that allows to match this target. To fix ideas, the parameter values for the whole period (1st column of Table 3) are:

$$\begin{array}{ccccccc} \alpha & \beta & \phi & \rho & \sigma & \delta & c \\ 0.675 & 0.951 & -0.107 & 0.936 & 0.183 & 0.044 & 0.581 \end{array}$$

In Table 3 and the subsequent tables, the confidence interval for mobility costs are constructed by evaluating the model at the lower and upper bounds of the confidence interval for σ (since, on the other hand, ϕ and ρ change little across periods). Mobility costs, which include one period of foregone earnings, are expressed as a fraction of average annual earnings.

The first remark concerns the levels of mobility costs reported in Table 3. We find that these costs fluctuate between 54 and 67 percent of annual earnings, depending on the period considered. We noted above that the *levels* of c should be interpreted with caution. Yet, they seem to “ring true” compared to estimates reported in the literature. For example, Kambourov and Manovskii (2009) observe that in the Handbook chapter of Heckman, LaLonde and Smith (1999), the average vocational training program has a direct cost of two months of wages for the median worker and takes about three months of full-time work, thus implying a total cost of five months of wages. Lee and Wolpin (2006) estimate that intersectoral mobility costs amount to about 75 percent of average annual earnings over the 1968–2000 period for male workers. They report that mobility costs are lower when switching occupations within the same sector. Thus, mobility costs amounting to 50 to 70 percent of annual earnings are well within the range of these estimates.

Table 3. Worker reallocation and mobility costs: Baseline results

	1976-2015	By subperiod			
		1976-1985	1986-1995	1996-2005	2006-2015
Excess reallocation	14.6	14.6	14.1	14.7	14.9
Mobility cost	0.581 [0.566,0.597]	0.546 [0.506,0.585]	0.544 [0.506,0.594]	0.538 [0.501,0.588]	0.672 [0.631,0.723]

NOTES: The table reports the estimates of excess reallocation reproduced from Panel B of Table 1 and the mobility cost predicted by the model. The curvature parameter is $\alpha = 0.675$ and the parameters of the stochastic process are those reported in Table 2. Figures in brackets are 95 percent confidence intervals.

The main result in Table 3 relates to changes in mobility costs across periods. To rationalize the behavior of excess worker reallocation over the 1976–2015 period, we find that mobility costs must remain steady during the first three decades, and that they must be substantially higher in the period 2006-2015. The stability in the costs of switching occupations during the period 1976-2005 is quite remarkable in light of the trends in reallocation that characterize the two decades 1986-1995 and 1996-2005. That is, according to the model, the decrease in excess reallocation followed by its increase is the mirror image of net reallocation captured by the parameter δ , so that productivity shocks and mobility costs play almost no role in these dynamics. More specifically, during the period from 1996

to 2005, the fact that δ decreases while σ experiences a slow-moving increase accounts for the higher levels of excess reallocation during that period.

Next, we remark on the results for the period 2006-2015, which illustrates well the need for a model. During the last decade, net (resp. excess) reallocation decreased (resp. increased) only slightly. Meanwhile, the volatility of productivity shocks continued its increase, so much so that the rates of excess reallocation could have been much higher than those observed in the data. The model attributes the rather steady levels of excess reallocation to a larger cost of switching occupation during this period. In terms of magnitude, this appears to be a substantial change: the mobility cost c increases by 25 percent relative to the period immediately before.

Discussion. The island model that we operationalize in this paper maps a stochastic process of productivity shocks, an exogenous reallocation event, and endogenous mobility decisions into a predicted level of excess reallocation. That model needs an additional parameter, the mobility cost c , in order to match the actual level of excess reallocation. Thus, one way to interpret c is that it is a residual that closes the gap between the model and data. In this respect, it is remarkable that the performance of the model seems to remain roughly constant during three decades. Notice, meanwhile, that the residual wages in equation (13) is allowed to capture a different amount of heterogeneity in different periods. This may lead to more “cleansing” of the data in some periods to align it to the model.

Why does the residual parameter c increase so much during the period that includes the Great Recession? The model is too stylized to provide an explanation, but we can at least speculate on the factors that might have prompted a shift towards higher mobility costs.

One line of interpretation is that mobility costs measure the loss of human capital that has been accumulated in a specific occupation or industry, as in, e.g., Rogerson (2005) or Kambourov and Manovskii (2009). In the model, mobility costs are homogeneous, which assumes implicitly that islands do not differ with respect to the amount of specific human capital that workers accumulate. There is ample evidence to argue against that assumption. To take one example, Meyer and Osborne (2005) show that there is a large dispersion in the amount of required vocational training across the 3-digit categories of the OCC1990 classification used in this paper.²⁴ Thus, we think that the increase in mobility costs could indicate that workers in jobs with a larger occupation- or industry-specific component were prompted to switch jobs during the period 2006-2015.

Another important line of interpretation is that mobility costs represent labor market frictions that prevent the right worker to be assigned readily to the right job. In this respect, we note that several studies have found an increase in occupational mismatch during the Great Recession: see Şahin et al. (2014), Barnichon and Figura (2015) and Herz and Van Rens (2015). The model we use does not allow for mismatch since workers can direct their search across islands. However, notice that workers must incur the mobility cost c to enter this directed search process. Thus, viewed through the lens of the model, the period surrounding the Great Recession is characterized by an increase in the cost paid to move to the most productive occupations.

²⁴For example, for “Retail sales clerks”, “Mail carriers for postal service” or “Paper folding machine operators”, the amount of special vocational training is estimated to be between 1 and 3 months. For “Physicians”, “Aerospace engineers” or “Lawyers”, the amount of training is estimated to be close to 10 years.

5.3. A Look at Reallocation Across Industries. In this subsection, we show that a similar picture emerges when the analysis is applied to worker reallocation across industries. We refer the reader to Appendix A.2 for the complete set of results. Briefly, for 3-digit industries, we find a number of now familiar patterns such as the long-run stability of the time series, the opposing trends during the periods 1986-1995 and 1996-2005, and the dip in excess reallocation during the Great Recession.

Table 4. Worker reallocation across industries and mobility costs

	1976-2015	By subperiod			
		1976-1985	1986-1995	1996-2005	2006-2015
Excess reallocation	9.74	9.42	9.23	9.96	10.18
Mobility cost	0.714 [0.676,0.739]	0.488 [0.446,0.537]	0.748 [0.676,0.810]	0.732 [0.671,0.796]	0.864 [0.801,0.951]

NOTES: The table reports the estimates of excess reallocation across industries and the mobility cost predicted by the model. The curvature parameter is $\alpha = 0.675$ and the parameters of the stochastic process are those reported in Table A4 in Appendix A. Figures in brackets are 95 percent confidence intervals.

Table A4 in the appendix displays the estimates of the parameters ϕ , ρ , σ using residual wages at the industry level.²⁵ The first column of Table 4 reports the value of mobility costs implied by excess reallocation across 3-digit industries. c amounts to 70 percent of annual earnings. Though this value is higher than the estimated costs of switching occupations, it is also more volatile across periods. Turning to the evolution over time, the change in the volatility of productivity shocks is broadly in line with that observed at the occupational level (Table 2 vs. Table A4). The timing of this change is different, however: at the industry level, we observe a steady increase in the parameter σ throughout the four decades. As shown in Table 4, there is also an upward trend in excess reallocation during the last two decades of the sample period. The main finding is that, through the lens of the model, the increase in excess reallocation would have been much larger had mobility costs not increased over time, and especially so during the last decade. This underscores one of our conclusions, that the increase in the volatility of productivity shocks and mobility costs may have been felt at a fine level of disaggregation, such as occupation-industry cells.

5.4. Sensitivity Analysis. In this subsection, we report the results from several numerical experiments which test the robustness of the main results. These sensitivity checks provide additional insights into the relationships between productivity shocks, mobility costs and worker reallocation.

Occupational Wages. In the first set of sensitivity checks, we study the effects of using a different measurement of occupational wages. To cover the spectrum of possible choices, we consider the 25th, 50th and 75th percentiles of the residual log-wage.²⁶ We use these alternative measurements together with the benchmark parameter $\alpha = 0.675$ to repeat the two steps of the estimation protocol. The

²⁵The specification of the residual-wage equation is the same as equation (13) except that industries dummies are used in lieu of occupation dummies.

²⁶We have also experimented with the 10th and 90th percentile of the residual log-wage. The results were similar but the parameters were less precisely estimated. Our preference for moments that are closer to the mean of the residual

complete set of estimates of the productivity process are reported in Table A1 in the Appendix. To focus attention on the key parameters, in Table 5 we report only the estimates of σ and the mobility cost c predicted by the model.

Relative to the mean residual log-wage, the 25th, 50th and 75th percentiles imply a stochastic process with a lower mean, less persistence and more volatility. The volatility is higher in the estimates delivered by the 25th and 75th percentiles relative to the 50th percentile, which suggests a U-shaped relationship between this parameter and the percentile of the wage distribution used to measure occupational wages.²⁷ Finally, and importantly, the three different measurements concur in delivering a higher volatility of productivity shocks during the period 2006-2015.

Since productivity shocks are more volatile when using the 25th, 50th and 75th percentiles of the residual log-wage, the mobility costs predicted by the model are also higher. This dovetails with our note of caution about interpreting the levels of the parameter c . Meanwhile, the levels of mobility costs displayed in Table 5 are not too far off compared to those in the benchmark results. The main finding is that the different panels of the table indicate: (i) mild upward changes in the parameters that govern excess worker reallocation during the years 1976-2005 and (ii) more volatile shocks and higher mobility costs during the period 2006-2015.²⁸

Curvature Parameter. In the second set of sensitivity checks, we revert to our benchmark definition of occupational wages (the mean residual log-wage) and analyze the effects of symmetric deviations of the curvature parameter around the value $\alpha = 0.675$. To remain within the range of plausible values for decreasing returns to scale in labor, we consider $\alpha = 0.600$ and $\alpha = 0.750$.

Table 6 reports the mobility costs obtained after changing the curvature parameter (the benchmark is reported in panel B of the table to facilitate comparisons). The results of the first step of the estimation procedure are summarized in Table A2: we find little differences relative to the benchmark and therefore we relegate the full results to the Appendix. The similarities can actually be ascertained by comparing mobility costs across panels in Table 6. Indeed, the levels shown in panels A and C are within the range of the baseline mobility costs. More curvature in the production function entails more persistence in shocks, which raises the unconditional mean of the productivity process in equation (9). This increases output and hence wages, so that mobility costs as a fraction of earnings are lower in Panel A. Finally, in this robustness check too, we find support for the main predictions of the model: mobility costs have remained steady during three decades while there has been a slight increase in the volatility of productivity shocks since the mid-1980s, and a more pronounced increase during the final period masked by mobility costs that were substantially higher.

Other Sensitivity Checks. We examined the effects of combining the robustness checks just described: we used different definitions of occupational wages (25th, 50th, 75th percentiles) combined with

log-wage distribution is based on the idea that the average/median labor market situation is more likely to act as a signal used by workers to direct their search away from or towards an occupation.

²⁷In fact, this U-shaped relationship is obscured by the inclusion of occupation dummies in the estimation of equation (13). The results without the occupation dummies are available upon request.

²⁸In Table 5, the cost of switching occupations seems to decrease during the 1996-2005 decade. This pattern is consistent with the increase in the relative share of *voluntary* labor market transitions during that period, which is documented in Borowczyk-Martins and Lalé (2016).

Table 5. Productivity shocks and mobility costs: Alternative occupational wages

	1976-2015	By subperiod			
		1976-1985	1986-1995	1996-2005	2006-2015
A. 25th percentile wage					
$\hat{\sigma}$	0.214	0.207	0.211	0.207	0.231
	[0.211,0.218]	[0.201,0.213]	[0.205,0.217]	[0.201,0.213]	[0.225,0.238]
\hat{c}	0.763	0.711	0.758	0.702	0.871
	[0.736,0.791]	[0.659,0.754]	[0.717,0.813]	[0.661,0.756]	[0.816,0.925]
B. 50th percentile wage					
$\hat{\sigma}$	0.212	0.206	0.206	0.204	0.231
	[0.209,0.215]	[0.200,0.212]	[0.200,0.212]	[0.198,0.210]	[0.224,0.237]
\hat{c}	0.715	0.709	0.725	0.648	0.835
	[0.701,0.764]	[0.657,0.750]	[0.674,0.776]	[0.598,0.698]	[0.783,0.900]
C. 75th percentile wage					
$\hat{\sigma}$	0.226	0.223	0.214	0.220	0.247
	[0.223,0.230]	[0.216,0.229]	[0.208,0.221]	[0.213,0.226]	[0.239,0.254]
\hat{c}	0.834	0.818	0.776	0.775	0.989
	[0.820,0.871]	[0.758,0.869]	[0.727,0.837]	[0.714,0.824]	[0.937,1.052]

NOTES: The table reports the estimates of the standard deviation of shocks and mobility cost predicted by the model, using alternative moment of the (residual log-) wage in each occupation. Panel A: 25th percentile; Panel B: 50th percentile; Panel C: 75th percentile. All estimations are performed with $\alpha = 0.675$. Figures in brackets are 95 percent confidence intervals.

Table 6. Reallocation and mobility costs: Alternative curvature parameters

	1976-2015	By subperiod			
		1976-1985	1986-1995	1996-2005	2006-2015
A. $\alpha = 0.600$					
\hat{c}	0.525	0.498	0.485	0.487	0.634
	[0.507,0.556]	[0.450,0.534]	[0.449,0.533]	[0.449,0.524]	[0.582,0.685]
B. $\alpha = 0.675$					
\hat{c}	0.581	0.546	0.544	0.538	0.672
	[0.566,0.597]	[0.506,0.585]	[0.506,0.594]	[0.501,0.588]	[0.631,0.723]
C. $\alpha = 750$					
\hat{c}	0.643	0.586	0.620	0.598	0.767
	[0.621,0.655]	[0.553,0.634]	[0.573,0.666]	[0.552,0.646]	[0.709,0.804]

NOTES: The table reports the mobility cost predicted by the model using the mean (residual log-) wage in each occupation and alternative values for the curvature parameter α . Panel A: $\alpha = 0.600$; Panel B: $\alpha = 0.675$; Panel C: $\alpha = 0.750$. Panel B reproduces Table 2 to facilitate comparisons. Figures in brackets are 95 percent confidence intervals.

different values for the curvature parameter ($\alpha = 0.600$, $\alpha = 0.675$, $\alpha = 0.750$). The changes that we obtained were minimal. The main reason is that, although the first step of the estimation procedure is contingent on the choice of α , the parameters of the productivity process turn out to be rather insensitive to this choice. As a result, the value for the standard deviation of shocks, σ , changes little across specifications, and so does the mobility cost c in the second step of the procedure.

Finally, using a different interaction between individual characteristics in the wage regressions (13) or including more controls gave similar results. That is, the parameters of the productivity process and mobility costs shift in levels, but their changes across periods deliver the same message as the benchmark estimates. These results also hold true when the sensitivity checks are applied to reallocation across industries. We conclude that the evolution of productivity shocks and mobility costs uncovered in this section is a robust prediction of the model.

6. CONCLUSION

We documented the evolution of worker reallocation across occupations in the U.S. labor market over a four- to five-decade period. Our empirical results complement and update existing findings on the extent of net reallocation and excess reallocation, and permit a broader view of the trends that have affected these time series. We went beyond the descriptive analysis: we used an equilibrium model to uncover potential changes in the factors that govern worker reallocation. Our findings indicate that the apparent long-run stability of reallocation across occupation is the result of slow-moving changes in the volatility of productivity shocks and mobility costs. In the recent period that includes the Great Recession, we find an upward shift in these variables: shocks were more volatile and this was accompanied by a marked increase in mobility costs. Viewed through the lens of the model, what characterizes recent years is an increase in the costs borne by workers to land the “right” occupations.

These findings are potentially relevant to several active research areas, including investigations of changes in the structure of the U.S. labor market and the job stability and security debates.

First of all, the period examined witnessed dramatic changes to the macro-environment – globalization, technological change, outsourcing or changes in labor force unionization, to name a few. In this respect, it is perhaps not surprising that one key parameter that drives worker reallocation, the volatility of productivity shocks, evolved in ways consistent with the view that the labor market becomes more turbulent. Our analysis also indicates that the cost of switching occupations did not decrease during that period, which dovetails well with job polarization, a leading paradigm to characterize labor market changes. In a somewhat stylized way, the analysis captures the fact that workers may incur higher mobility costs because declining islands – jobs involving routine tasks – and expanding islands – non-routine cognitive and non-routine manual jobs – are drifting away from each other.

Second, the finding that mobility costs increased in recent years is also consistent with several trends in the U.S. labor market which relate to job stability and job security. For instance, it is well known that job-to-job transitions have become less frequent since the late 1990s. To the extent that many job changes are accompanied by a change in occupation, the increase in mobility cost that we uncover could potentially contribute to this trend. Another example is the contemporaneous decrease in the number of single jobholders who take on a second job. Since workers usually moonlight in

an occupation that differs from that of the primary job, this trend could also be linked to the upward shift in mobility costs. From a policy standpoint, understanding whether the apparent decline in turnover reflects increased job stability or fewer opportunities for workers to change their career paths is crucial, and should be an avenue for future research.

REFERENCES

- Abraham, Katharine G, and Lawrence F Katz.** 1986. "Cyclical unemployment: Sectoral shifts or aggregate disturbances?" *Journal of Political Economy*, 94(3): 507–522.
- Alvarez, Fernando, and Marcelo Veracierto.** 2000. "Labor-market policies in an equilibrium search model." In *NBER Macroeconomics Annual*. Vol. 14, 265–316. MIT.
- Alvarez, Fernando, and Marcelo Veracierto.** 2012. "Fixed-term employment contracts in an equilibrium search model." *Journal of Economic Theory*, 147(5): 1725–1753.
- Alvarez, Fernando, and Robert Shimer.** 2011. "Search and rest unemployment." *Econometrica*, 79(1): 75–122.
- Barnichon, Regis, and Andrew Figura.** 2015. "Labor market heterogeneity and the aggregate matching function." *American Economic Journal: Macroeconomics*, 7(4): 222–249.
- Borowczyk-Martins, Daniel, and Etienne Lalé.** 2016. "The rise of part-time employment." *Sciences Po Discussion paper 2016-04*.
- Carrillo-Tudela, Carlos, and Ludo Visschers.** 2013. "Unemployment and endogenous reallocation over the business cycle." *CESifo Working Paper 4079*.
- Coen-Pirani, Daniele.** 2010. "Understanding gross worker flows across US states." *Journal of Monetary Economics*, 57(7): 769–784.
- Davis, Morris A, Jonas DM Fisher, and Marcelo Veracierto.** 2016. "Gross migration, housing and urban population dynamics." *Mimeo, previously FRB of Chicago Working Paper 2013-19*.
- Davis, Steven J.** 2008. "The decline of job loss and why it matters." *American Economic Review: Papers and Proceedings*, 98(2): 263–267.
- Davis, Steven J, R Jason Faberman, and John Haltiwanger.** 2006. "The flow approach to labor markets: New data sources and micromacro links." *Journal of Economic Perspectives*, 20(3): 3–26.
- den Haan, Wouter J, Christian Haefke, and Garey Ramey.** 2005. "Turbulence and unemployment in a job matching model." *Journal of the European Economic Association*, 3(6): 1360–1385.
- Fujita, Shigeru.** 2015. "Declining labor turnover and turbulence." *FRB of Philadelphia Working Paper No. 15-29*.
- Gallipoli, Giovanni, and Gianluigi Pelloni.** 2014. "Macroeconomic effects of job reallocations: A survey." *Review of Economic Analysis*, 5(2): 127–176.
- Heckman, James J, Robert J LaLonde, and Jeffrey A Smith.** 1999. "The economics and econometrics of active labor market programs." *Handbook of Labor Economics*, 3: 1865–2097.
- Herz, Benedikt, and Thijs Van Rens.** 2015. "Accounting for mismatch unemployment." *IZA Discussion Paper No. 8884*.
- Hyatt, Henry R, and James R Spletzer.** 2013. "The recent decline in employment dynamics." *IZA Journal of Labor Economics*, 2(1): 1–21.

- Jovanovic, Boyan, and Robert Moffitt.** 1990. "An estimate of a sectoral model of labor mobility." *Journal of Political Economy*, 98(4): 827–852.
- Kambourov, Gueorgui, and Iourii Manovskii.** 2008. "Rising occupational and industry mobility in the United States: 1968-1997." *International Economic Review*, 49(1): 41–79.
- Kambourov, Gueorgui, and Iourii Manovskii.** 2009. "Occupational mobility and wage inequality." *Review of Economic Studies*, 76(2): 731–759.
- Kambourov, Gueorgui, and Iourii Manovskii.** 2013. "A cautionary note on using (March) CPS and PSID data to study worker mobility." *Macroeconomic Dynamics*, 17(01): 172–194.
- King, Miriam, Steven Ruggles, J. Trent Alexander, Sarah Flood, Katie Genadek, Matthew B. Schroeder, Brandon Trampe, and Rebecca Vick.** 2010. "Integrated public use microdata series, Current Population Survey: Version 3.0. [Machine-readable database]." *Minneapolis: University of Minnesota*.
- Lalé, Etienne.** 2012. "Trends in occupational mobility in France: 1982–2009." *Labour Economics*, 19(3): 373–387.
- Lalé, Etienne.** 2016. "Turbulence and the employment experience of older workers." *IZA Discussion Paper No. 10061*.
- Layard, R., P.R.G. Layard, S. Nickell, and R. Jackman.** 2005. *Unemployment: Macroeconomic performance and the labour market*. Oxford University Press, USA.
- Lee, Donghoon, and Kenneth I Wolpin.** 2006. "Intersectoral labor mobility and the growth of the service sector." *Econometrica*, 74(1): 1–46.
- Lilien, David M.** 1982. "Sectoral shifts and cyclical unemployment." *Journal of Political Economy*, 90(4): 777–793.
- Ljungqvist, Lars, and Thomas J Sargent.** 1998. "The European unemployment dilemma." *Journal of Political Economy*, 106(3): 514–550.
- Ljungqvist, Lars, and Thomas J Sargent.** 2008. "Two questions about European unemployment." *Econometrica*, 76(1): 1–29.
- Lkhagvasuren, Damba.** 2012. "Big locational unemployment differences despite high labor mobility." *Journal of Monetary Economics*, 59(8): 798–814.
- Lucas, Robert E, and Edward C Prescott.** 1974. "Equilibrium search and unemployment." *Journal of Economic Theory*, 7(2): 188–209.
- Meyer, Peter B, and Anastasiya M Osborne.** 2005. "Proposed category system for 1960-2000 Census occupations." *US Bureau of Labor Statistics, Working Papers Series*, 383.
- Moscarini, Giuseppe.** 2001. "Excess worker reallocation." *Review of Economic Studies*, 68(3): 593–612.
- Moscarini, Giuseppe, and Francis G Vella.** 2003. "Aggregate worker reallocation and occupational mobility in the United States: 1976-2000." *The Institute for Fiscal Studies Working Paper*.
- Moscarini, Giuseppe, and Francis G Vella.** 2008. "Occupational mobility and the business cycle." *NBER Working Paper No. 13819*.
- Moscarini, Giuseppe, and Kaj Thomsson.** 2007. "Occupational and job mobility in the US." *Scandinavian Journal of Economics*, 109(4): 807–836.

- Murphy, Kevin M, and Robert H Topel.** 1987. “The evolution of unemployment in the United States: 1968-1985.” *NBER Macroeconomics Annual*, 2: 11–58.
- Phelps, Edmund S, Armen A Alchian, Charles C Holt, and Dale T Mortensen.** 1970. *Microeconomic foundations of employment and inflation theory*. W. W. Norton and Company.
- Rogerson, Richard.** 2005. “Sectoral shocks, human capital, and displaced workers.” *Review of Economic Dynamics*, 8(1): 89–105.
- Şahin, Ayşegül, Joseph Song, Giorgio Topa, and Giovanni L Violante.** 2014. “Mismatch unemployment.” *American Economic Review*, 104(11): 3529–3564.
- Tauchen, George.** 1986. “Finite state Markov-chain approximations to univariate and vector autoregressions.” *Economics Letters*, 20(2): 177–181.
- U.S. Bureau of the Census.** 2010. “Current Population Survey - Displaced worker, employee tenure, and occupational mobility supplement file.” *Technical documentation available at: <https://www.census.gov/cps/>*.
- Van Nieuwerburgh, Stijn, and Pierre-Olivier Weill.** 2010. “Why has house price dispersion gone up?” *Review of Economic Studies*, 77(4): 1567–1606.
- Wiczer, David G.** 2015. “Long-term unemployment: Attached and mismatched?” *Mimeo, Federal Reserve Bank of St. Louis*.

APPENDICES

APPENDIX A. ADDITIONAL RESULTS

In Subsection A.1 in this appendix, we report the complete set of estimates of the productivity process used in the numerical experiments. In Subsection A.2, we provide the complete results for worker reallocation across industries.

A.1. Estimates of the productivity process. Table A1 reports the estimates of the productivity process obtained using the 25th, 50th and 75th percentiles of the residual log-wage to measure occupational wages. The curvature parameter is the same as under the benchmark, i.e. $\alpha = 0.675$.

Table A1. Estimates of the productivity process: Alternative occupational wages

	1976-2015	By subperiod			
		1976-1985	1986-1995	1996-2005	2006-2015
A. 25th percentile wage					
$\hat{\phi}$	-0.170	-0.166	-0.169	-0.165	-0.182
	[-0.187,-0.153]	[-0.198,-0.134]	[-0.202,-0.136]	[-0.199,-0.132]	[-0.218,-0.145]
$\hat{\sigma}$	0.214	0.207	0.211	0.207	0.231
	[0.211,0.218]	[0.201,0.213]	[0.205,0.217]	[0.201,0.213]	[0.225,0.238]
$\hat{\rho}$	0.914	0.916	0.914	0.916	0.910
	[0.906,0.922]	[0.900,0.931]	[0.897,0.930]	[0.900,0.933]	[0.893,0.928]
B. 50th percentile wage					
$\hat{\phi}$	-0.137	-0.133	-0.133	-0.141	-0.141
	[-0.151,-0.123]	[-0.161,-0.106]	[-0.160,-0.105]	[-0.169,-0.113]	[-0.171,-0.111]
$\hat{\sigma}$	0.209	0.205	0.204	0.200	0.225
	[0.206,0.212]	[0.199,0.211]	[0.198,0.210]	[0.195,0.206]	[0.218,0.231]
$\hat{\rho}$	0.918	0.920	0.921	0.914	0.918
	[0.910,0.926]	[0.904,0.935]	[0.905,0.937]	[0.898,0.930]	[0.901,0.935]
C. 75th percentile wage					
$\hat{\phi}$	-0.128	-0.129	-0.122	-0.131	-0.132
	[-0.141,-0.115]	[-0.155,-0.103]	[-0.147,-0.097]	[-0.156,-0.106]	[-0.159,-0.104]
$\hat{\sigma}$	0.226	0.223	0.214	0.220	0.247
	[0.223,0.230]	[0.216,0.229]	[0.208,0.221]	[0.213,0.226]	[0.239,0.254]
$\hat{\rho}$	0.907	0.907	0.911	0.904	0.906
	[0.898,0.915]	[0.889,0.924]	[0.894,0.928]	[0.887,0.921]	[0.888,0.925]

NOTES: The table reports the estimates of the productivity process obtained using alternative moment of the (residual log-) wage in each occupation. Panel A: 25th percentile; Panel B: 50th percentile; Panel C: 75th percentile. All estimations are performed with $\alpha = 0.675$. Figures in brackets are 95 percent confidence intervals.

Table A2 reports the estimates of the productivity process that we obtain for different values of the curvature parameter α . The definition of occupational wages is the same as under the benchmark, i.e. we use the mean residual log-wage.

Table A2. Estimates of the productivity process: Alternative curvature parameters

1976-2015	By subperiod				
	1976-1985	1986-1995	1996-2005	2006-2015	
A. $\alpha = 0.600$					
$\hat{\phi}$	-0.111	-0.114	-0.108	-0.113	-0.108
	[-0.125,-0.096]	[-0.142,-0.087]	[-0.135,-0.080]	[-0.142,-0.085]	[-0.138,-0.078]
$\hat{\sigma}$	0.201	0.197	0.192	0.196	0.217
	[0.198,0.203]	[0.191,0.202]	[0.186,0.197]	[0.190,0.202]	[0.210,0.223]
$\hat{\rho}$	0.946	0.943	0.947	0.944	0.949
	[0.939,0.952]	[0.930,0.956]	[0.934,0.960]	[0.930,0.957]	[0.935,0.963]
B. $\alpha = 0.675$					
$\hat{\phi}$	-0.107	-0.108	-0.103	-0.110	-0.110
	[-0.120,-0.095]	[-0.132,-0.083]	[-0.127,-0.078]	[-0.135,-0.084]	[-0.137,-0.083]
$\hat{\sigma}$	0.183	0.178	0.175	0.179	0.200
	[0.181,0.186]	[0.173,0.183]	[0.170,0.180]	[0.174,0.184]	[0.194,0.205]
$\hat{\rho}$	0.936	0.935	0.939	0.934	0.936
	[0.929,0.943]	[0.921,0.949]	[0.925,0.953]	[0.919,0.948]	[0.921,0.952]
C. $\alpha = 0.750$					
$\hat{\phi}$	-0.104	-0.101	-0.097	-0.106	-0.113
	[-0.115,-0.093]	[-0.122,-0.080]	[-0.119,-0.076]	[-0.128,-0.084]	[-0.137,-0.089]
$\hat{\sigma}$	0.168	0.161	0.161	0.164	0.185
	[0.166,0.170]	[0.157,0.166]	[0.156,0.165]	[0.159,0.169]	[0.180,0.190]
$\hat{\rho}$	0.920	0.922	0.925	0.918	0.915
	[0.912,0.928]	[0.906,0.937]	[0.910,0.941]	[0.902,0.934]	[0.898,0.933]

NOTES: The table reports the estimates of the productivity process obtained using the mean (residual log-) wage in each occupation and alternative values for the curvature parameter α . Panel A: $\alpha = 0.600$; Panel B: $\alpha = 0.675$; Panel C: $\alpha = 0.750$. Panel B reproduces Table 2 to facilitate comparisons. Figures in brackets are 95 percent confidence intervals.

A.2. Worker Reallocation across Industries. The charts in Figure A1 display net reallocation at the different digit levels of the industry classification. The fraction of total employment which is reallocated across 3-digit industries is on average 3.9 percent per year. As can be observed, the time series exhibits an upward trend in the 1980s and early 1990s, and then it remains stable until the end of the sample period. The increase between 1986 and 1995 is statistically significant (Table A3). Finally, similar to net reallocation across occupations, there is no apparent relationship between net reallocation across industries and the recessionary periods covered by the data.

Figure A2 shows the evolution of excess reallocation across industries since 1980. Excess worker reallocation is on average 9.74 percent at the 3-digit level, which is much lower than excess reallocation across occupations. Meanwhile, as noted in the text, its contribution to overall reallocation (the sum of net and excess reallocation) is similar to what we obtained for occupations, i.e. a contribution of about 70 percent. Next, at the 3-digit level we notice some patterns that are similar to those in Figure 2: excess reallocation has been stable in the 1980s and early 1990s, (ii) it has been on an upward

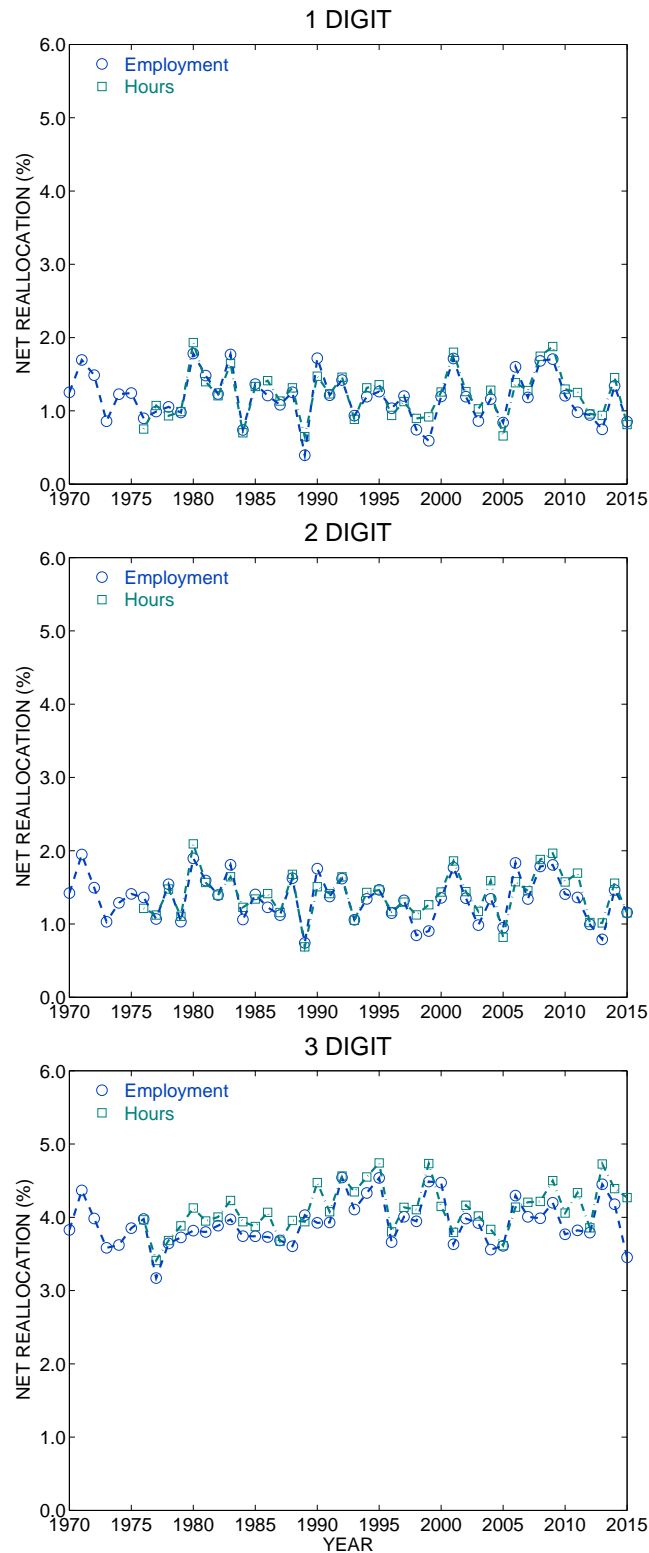


Figure A1. Net reallocation across industries

The upper, middle and lower charts display, respectively, net reallocation rates at the 1-, 2- and 3-digit level of the industry classification. Circles and squares denote, respectively, employment-weighted and hours-weighted rates of net reallocation. The hours variable is not available prior to 1976, and so the hours-weighted time series are only from 1976 onwards.

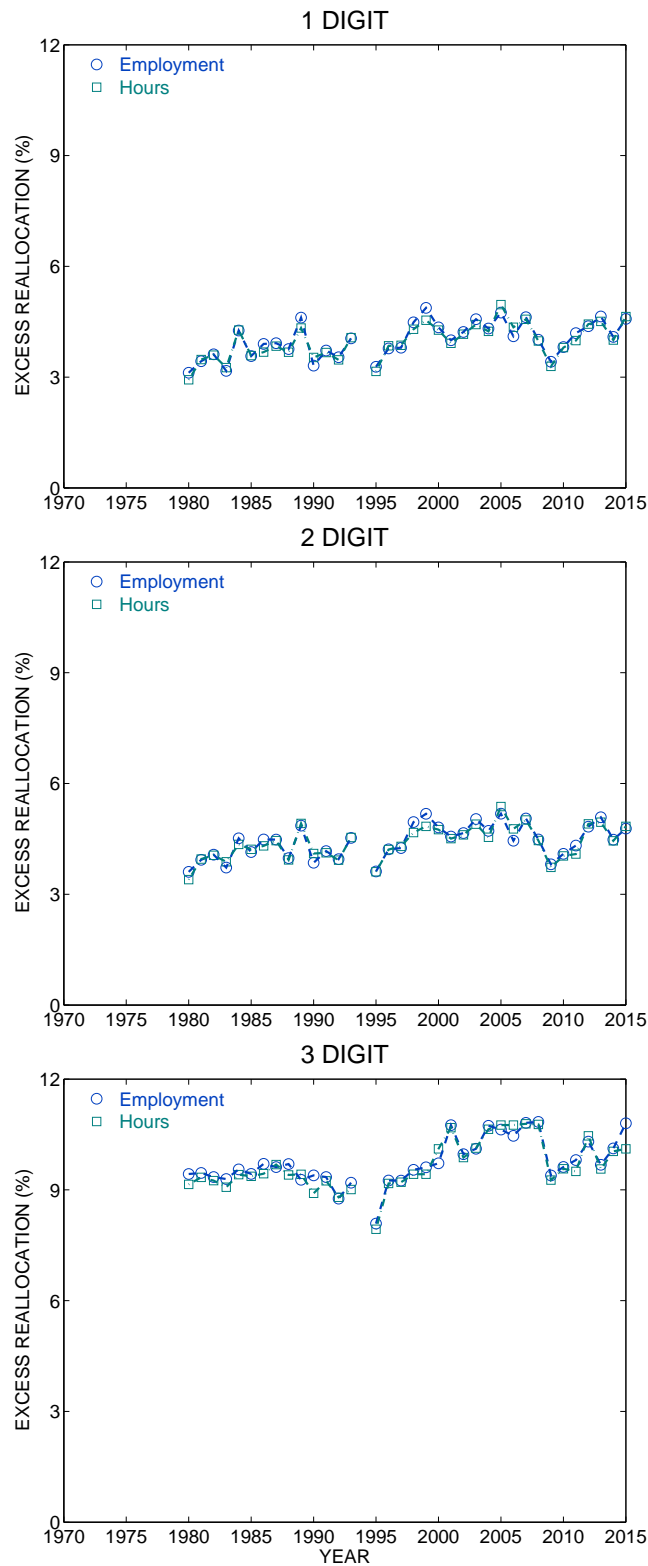


Figure A2. Excess reallocation across industries

The upper, middle and lower charts display, respectively, excess reallocation rates at the 1-, 2- and 3-digit level of the industry classification. Circles and squares denote, respectively, employment-weighted and hours-weighted rates of excess reallocation. The time series begin in 1980 which is the first period of observation for gross industry mobility (see Appendix B.2).

course since the mid-1990s, and (iii) it has dropped during the Great Recession. At the 1-digit and 2-digit levels, generally there has been an increase in excess reallocation across industries during the whole sample period.

Table A3. Worker reallocation across industries: Levels and trends

	1976-2015	By subperiod			
		1976-1985	1986-1995	1996-2005	2006-2015
A. Net reallocation					
Level	3.93 (0.05)	3.75 (0.07)	4.04 (0.11)	3.93 (0.11)	3.99 (0.09)
Trend	0.67 (0.42)	2.41 (2.54)	9.60 (2.01)	-3.70 (3.70)	-3.18 (3.27)
B. Excess reallocation					
Level	9.74 (0.11)	9.42 (0.04)	9.23 (0.17)	9.96 (0.18)	10.18 (0.17)
Trend	3.31 (0.84)	0.71 (2.39)	-15.52 (3.24)	16.82 (3.31)	-2.82 (6.27)

NOTES: Net reallocation and excess reallocation across 3-digit industries measured using the employment weights (circles, lower chart in Figures A1 and A2). The trend is computed using a linear regression of the time-series against calendar years; the coefficient is multiplied by 100 for legibility. Standard errors in parentheses.

Table A3 summarizes the evolution of net reallocation and excess reallocation across industries at the 3-digit level. These levels of worker reallocation are then used to inform the model. The estimates of the productivity process that we rely on for that exercise are displayed in Table A4. They are based on industry wages which we set to the mean residual log-wage in each industry. The curvature parameter is: $\alpha = 0.675$. The results of the numerical exercise are analyzed in Subsection 5.3.

Table A4. Estimates of the productivity process at the industry level

	1976-2015	By subperiod			
		1976-1985	1986-1995	1996-2005	2006-2015
$\hat{\phi}$	-0.071 [-0.083,-0.059]	-0.047 [-0.067,-0.026]	-0.078 [-0.102,-0.054]	-0.078 [-0.102,-0.053]	-0.078 [-0.104,-0.053]
$\hat{\sigma}$	0.137 [0.135,0.140]	0.111 [0.108,0.115]	0.135 [0.130,0.139]	0.143 [0.138,0.148]	0.156 [0.151,0.161]
$\hat{\rho}$	0.956 [0.949,0.963]	0.973 [0.960,0.986]	0.951 [0.936,0.966]	0.952 [0.937,0.967]	0.952 [0.936,0.967]

NOTES: The table reports the estimates of the productivity process obtained using the mean (residual log-) wage in each industry. The curvature parameter is $\alpha = 0.675$. Figures in brackets are 95 percent confidence intervals.

APPENDIX B. MEASUREMENT ISSUES

An article by Kambourov and Manovskii (2013) warns against certain pitfalls of using the March CPS to study worker reallocation across occupations. In Subsection B.1, we summarize their arguments and explain why the measurement of net reallocation is likely immune to most shortcomings that the authors identified. In Subsection B.2, we explain how we construct our measurement of excess reallocation using other data sources than the March CPS.

B.1. Net reallocation. Kambourov and Manovskii (2013) argue that the March CPS suffers from: (i) measurement error in the occupation of employment of respondents (ii) potential changes in the amount of noise in the final data due to changes in imputation techniques in 1976 and in 1989 and (iii) uncertainty about the time horizon over which mobility is measured, since the longest job held in the previous year does not necessarily coincide with the job held in March of the previous year.

The first issue – measurement error in occupational affiliation – is a hurdle faced by virtually any study of worker mobility. Meanwhile, there are reasons to believe that it is less of a problem for the measurement of net reallocation. Firstly, the occupational classification we use (the so-called OCC1990 classification) aggregates “close” occupational categories of the original CPS data. This has the potential of reducing the impact of coding error.²⁹ Second, when a worker whose “true” occupation of employment is A is erroneously classified in occupation B, it may well be that a worker in occupation B is misclassified in occupation A in the same CPS file. These errors have no impact on net reallocation because they do not affect occupational employment shares.

The second problem – changes in imputation techniques – is also less likely to matter for net flows the way it does for gross flows. Imputations techniques may generate spurious individual transitions across occupations as knowledge of both the current and past occupation is required to identify these transitions. In contrast, net reallocation relies only on cross-sectional information (one occupation per individual). Finally, we find no evidence of a break in the time series reported in the paper in 1976 or 1989, when the imputation techniques of the March CPS were changed.

The third problem – uncertainty about the time horizon between the current job and that which the individual describes as her/his main job in the previous year – is inconsequential for net reallocation. Indeed, to calculate net reallocation between time $t - 1$ and time t , we use the two cross sections of individuals surveyed at time $t - 1$ and time t , instead of drawing information about the past and present occupation of employment of individuals from the time- t cross section only.³⁰

In fact, we concur with Kambourov and Manovskii (2013) in suspecting that, in the March CPS files, the longest job held in the previous year is often different from the job held in March of the previous year. That is, we find that the rates of net reallocation computed using the past and present occupation of employment from the time- t cross section are systematically lower than those displayed in Figure 1. This is consistent with a shorter time horizon, which implies that there have been fewer changes in the occupational distribution of employment relative to changes over a one-year window.

²⁹For instance, “Statistical clerks” may be erroneously coded “Data-entry keyers” but are less likely to be erroneously coded as “Barbers”. The erroneous code in the original data vanishes away when “Statistical clerks” and “Data-entry keyers” are combined into a single occupation in the OCC1990 classification.

³⁰The time- t cross section of the March CPS has information about occupations held at time $t - 1$ and time t .

B.2. Excess Reallocation.

The MORG Files. To calculate gross worker flows (cf. equation (2)), we use the monthly files of the CPS Outgoing Rotation Group samples distributed by the National Bureau of Economic Research for the period from 1979 onwards. The so-called MORG files are extracts of the monthly CPS that correspond to a household's fourth and eighth month in the survey, which have a one-year gap between them. In these extracts, CPS respondents report their weekly earnings and hours in addition to the regular survey items. We use these information to implement sample restrictions that are very similar to those used for the measurement of net reallocation (see the online appendix).

We longitudinally match individuals using household and person identifiers combined with an age/sex/race filter. There have been changes in the CPS identifiers that prevents us from linking individuals in certain subperiods. In particular, as can be seen in Figure 2, it is impossible to obtain even one year-to-year match in 1994. Next, using matched individuals who are employed both during their fourth and eighth interview month, we compute annual gross flows across occupations and industries. Notice that, so doing, we obtain *monthly* time series of the *annual* gross flows. We convert these series to an annual frequency by taking the average of the monthly values.

One issue with these gross flow rates is that they are likely upward-biased by measurement error in occupational and industry affiliations. Moscarini and Thomsson (2007) show that (monthly) occupational mobility rates drop from 34 percent to less than 4 percent when spurious transitions are discarded from the 1979-1993 files of the CPS. At the annual frequency, we find that gross flows rates in the raw data average at 46.7 percent and 33.6 percent for 3-digit occupations and industries, respectively. This is substantially higher than our final estimates: the respective corresponding figures are 19.0 percent and 13.7 percent (details follow).

The Occupational Mobility Supplements. To overcome the issue of measurement error, we combine our MORG-based time series with estimates that we obtain from the Job Tenure and Occupational Mobility supplements of the CPS. These supplements have been administered every two years between 1996 and 2014 either in January or February. They contain information about an individual's occupation and industry from one year ago, and how long s/he has worked at her/his current job.

The Occupational Mobility supplements give only discrete snapshots of worker reallocation. This drawback is offset by the fact that they are likely to provide reliable estimates of the annual gross flows rates. One reason for this is that the interview technique uses dependent coding and refers explicitly to the period one year prior to the survey. To take one example from the January 2010 questionnaire, the respondent is asked to answer the following questions:

[ST20] Earlier you told me that (name/you) (is/are) now working as (FILL OCCUPATION FROM CPS). (Was/Were) (name/you) doing the same kind of work a year ago, in January 2009?

[ST21] What kind of work did (name/you) do, that is, what was (his/her/your) occupation in January 2009? (U.S. Bureau of the Census, 2010)

It has been established, in the context of occupational mobility, that dependent coding allows to eliminate a substantial amount of measurement error (Moscarini and Thomsson, 2007; Lalé, 2012).

Moreover, it is clear that the above survey questions alleviate uncertainty about the time horizon between the two occupations of employment of the respondent.

Thus, after obtaining the estimates of annual occupation and industry mobility rates from these CPS supplements, we use them to scale down the MORG-based time series. We multiply each of the time series by an adjustment factor that makes the average over the years 1996 to 2014 consistent with the average value from the Occupational Mobility supplements. We compute excess reallocation by subtracting net reallocation based on the March CPS from the adjusted gross flow rates.

APPENDIX C. COMPUTATIONAL DETAILS

Our numerical methodology to compute the equilibrium of the model is as follows:

- (1) We discretize the support of productivity and labor force sizes. For productivity, we use 35 grid points and apply standard approximation methods to compute the transition probabilities of the stochastic process. We use 3,500 grid points for the size of the labor force. Notice that the ergodic set of labor force sizes is an endogenous object.
- (2) We iterate until convergence on the Bellman equation that governs a worker's behavior. After convergence is obtained, we compute the employment policy function $n(z, \ell)$.
- (3) We recover the average size of the labor force implied by v_s by simulating the model. Specifically, we simulate the economy for 20,200 periods, the first 200 of which are discarded. We repeat this procedure 500 times, and the final labor force size is obtained by averaging over the 500 samples of 20,000 periods.
- (4) We update v_s accordingly, and repeat the different steps until convergence.