A Dynamic Model of Occupational Mobility, Structural Unemployment, Average Labour Productivity and Wage Dispersion

Fabrice Collard
The University of Bern

Jacob Wong*
The University of Adelaide

This version : March 24, 2011

Abstract

This paper presents a dynamic model of structural unemployment and occupational choice in which an economy is subjected to aggregate reallocation shocks and workers may choose to incur costs to retrain in order to move into occupations that pay higher wages. As it is costly for workers to retrain, workers may prefer to remain unemployed in occupations suffering through relatively low productivity states. Reallocation shocks, which change the relative labour productivity across occupations, drive variation in the distribution of workers across occupations. The general equilibrium model produces volatile unemployment and occupational mobility rates but also features (i) sign switches in the correlation between average labour productivity and unemployment rates across decades within sample runs, (ii) decades with negatively correlated unemployment and occupational mobility rates, and (iii) substantially large swings in the rate of occupational mobility that are strongly, positively correlated with measures of wage inequality.

*Corresponding author. School of Economics, The University of Adelaide, Ground Floor, Napier Building, Adelaide, SA 5005, Australia. Email : jacob.wong@adelaide.edu.au. We thank Chris Edmond, Viktoria Hnatkovska, Nir Jaimovich, Amartya Lahiri, Henry Siu and Lawrence Uren for valuable comments. We also thank seminar participants at the University of Melbourne, UNSW, the Reserve Bank of Australia and the Australian Conference on Quantitative Macroeconomics. Finally, we thank Chris Robinson and Giuseppe Moscarini for providing useful data.

This paper presents a dynamic model of occupational choice in which an economy is subjected to aggregate reallocative shocks and workers may choose to incur costs to retrain in order to move to higher wage occupations. We study how the individual incentives to remain "rest" unemployed or to retrain for new occupations interact with variation in the aggregate distribution of workers across occupations in order to yield changes in the correlation between the unemployment rate and average labour productivity over time. As the interaction between individual incentives to switch occupations and the distribution of workers across occupations is the primary mechanism at work in this model economy, we examine the implications of this interaction on the dynamics of the relationships between wage dispersion across occupations with the unemployment rate, the rate of occupation switching and average labour productivity.

It is shown that such a model can deliver time series which have three features which resemble those of the U.S. economy over the last few decades. First, average labour productivity and unemployment rates can exhibit negative contemporaneous correlations over lengthy periods and positive contemporaneous correlations over other periods. Second, the model economy typically generates a positive correlation between wage inequality and occupational mobility. Third, the model exhibits extended periods in which the rate of occupational mobility and the unemployment rate are negatively correlated. Importantly, the model economy is only subjected to shocks which change the relative labour productivity across occupations. We refer to this type of shock as a reallocation shock. It turns out that the history of the distribution of workers across occupations plays a primary role in generating the results.

In our model, workers are distributed across a continuum of occupations with each occupation producing an intermediate good that may be used in the production of a final consumption good. At any point in time a given worker possesses skills to work in one particular occupation. Every occupation has its own level of labour productivity and all occupations experience fluctuations in labour productivity over time. Each period, a given worker faces an idiosyncratic fixed cost of retraining. Furthermore, if a worker chooses to retrain, the cost of switching to an occupation that uses different tasks relative to the worker's current occupation is high in comparison to switching to an occupation using similar tasks relative to the worker's current occupation. Workers who are skilled in an occupation where productivity is sufficiently high to pay a participation wage face the option of working or not working for the period and retraining in order to have the skills to work in a different occupation next period. Workers who are skilled in an occupation where productivity is not high enough to pay a participation wage face the option of being rest unemployed for the period and remaining in the occupation next period versus retraining and having skills to work in a different occupation next period. Over time, the economy is buffeted by reallocative shocks which change the relative productivity across occupations. We interpret this as the effect of technological innovations that favour some occupations at the expense of others. The effect of the fixed costs of retraining as well as convex costs of retraining to dissimilar occupations is that the distribution of workers across occupations changes over time. This in turn affects the wage distribution across occupations. The distribution of wages feeds back into the equilibrium decisions of workers to retrain generating an interesting interplay between the distribution of employment across occupations, aggregate labour productivity and the unemployment rate.

Solving for the equilibrium of this economy requires the ability to keep track of the distributions of workers across occupations as well as the distribution of labour productivity across occupations. We make two simple assumptions in order to manage the computational difficulties of tracking these two infinite dimensional objects. The first assumption is that occupations are located on the circumference of a circle; in other words, there is an occupation ring. The greater the distance between two occupations on the ring, the more dissimilar the tasks used by the occupations in production. The second modeling assumption is that the relative productivities across occupations in relation to the most productive occupation is preserved across time even though the identity of the most productive occupation can change over time. In this specific sense, a reallocation shock can be modeled simply as a change in the identity of the most productive occupation. Once the identity of the most productive occupation is known the relative productivity of all other occupations can be determined. These modeling assumptions reduce the computational problem to that of tracking the distribution of workers across occupations and tracking the identity of the most productive occupation.

As an application of this model featuring aggregate fluctuations in the rate of occupation switching, we highlight two interesting properties of this model. Although the economy is buffeted only by reallocative shocks, following a shock, average labour productivity and the unemployment rate either move in the same direction, or they move in the opposite direction. Which situation arises depends crucially on the distribution of workers across occupations along with which occupations are favoured by the change in technology. If the change in technology just happens to favour an occupation in which a large fraction of workers are currently employed then there will be an increase in average productivity. Should it be that simultaneously, many workers are employed in disadvantaged occupations then the unemployment rate will rise. The opposite case can also occur, in which fewer workers are stuck in disadvantaged occupations so that the unemployment rate will fall. Generally, there is much history dependence in the qualitative response to a reallocation shock of a given size.

The other interesting property of this model is that we can track the distribution of wages across occupations over time and relate this to the dynamic behaviour of occupation switching and rest unemployment. We find that in periods of high occupational mobility, measures of wage inequality tend to rise. This result is intuitive given the mechanism at work in the economy. When the economy does not experience any large reallocation shocks for awhile, workers switch to higher paying occupations over time. This leads to workers bunching into occupations with high levels of labour productivity. If the economy is then hit by a reallocation shock, some occupations suddenly experience a reduction in labour productivity and these occupations yield lower wages. Some workers are then stuck in occupations with very low wages and choose to be rest unemployed. Not all workers who suffer wage losses choose to switch occupations immediately due to potentially high idiosyncratic fixed costs of occupation switching. This fraction of the population increases the measure of workers obtaining low wages. Others who suffer wage reductions will choose to incur costs in order to switch occupations and will move to occupations paying higher wages and this movement gives rise to the increase in wage inequality. During this transition we observe a rise in the unemployment rate, a rise in occupation switching and a rise in wage inequality.

As described above, in order to generate changes in the qualitative relationship between the unemployment rate and average labour productivity (and also changes in the relationship between the unemployment rate and occupational mobility rate) the economy requires changes in the rate of occupation switching over time. The model economy is able to generate large endogenous swings in the rate of occupational switching in simulations for which parameters are chosen such that median distance of individual moves across occupations as well as the average frequency of occupation switches resemble statistics on occupational changes as documented by Robinson (2010). As far as we know, this is the first paper in which there are aggregate fluctuations in the distribution of workers across a large number of occupations and that accounts for the individual's choice of the "distance" between current and past occupations.

The literature most related to the model in this paper is that which has grown from the seminal paper of equilibrium unemployment by Lucas and Prescott (1974). In their paper, Lucas and Prescott introduce a theory of equilibrium unemployment in which workers are attached to islands which, in our context, may be thought of as occupations. Occupations are subject to idiosyncratic demand shocks which affect the wages paid to workers attached to that occupation. Workers may choose to work at the prevailing wage, leave the occupation and search for new jobs or stay attached to occupations and wait for better times. The main difference with our model is that in our model there is a measure of distance between occupations and it is more costly for workers to move to occupations that use dissimilar tasks than the worker's current occupation. In the classic Lucas-Prescott island model, when workers decide to switch islands, they sit out of the labour force for a period and then have unrestricted choice of islands upon which to reenter the labour force. Alvarez and Shimer (2009a) and Alvarez and Shimer (2009b) present extensions of the Lucas-Prescott model and using continuous-time dynamic programming techniques are able to present elegant closed-form solutions to a version of the economy's steady state.[1] While both models provide insights into the theory of rest unemployment, and solve for the steady state in which workers direct their choice of occupation switches, neither paper examines

[1] It is worth noting that Alvarez and Shimer (2009a) use their model to understand wage behaviour at the industry level whereas our application is to highlight the dynamics of the correlations between the unemployment rate, average labour productivity, occupational mobility rates and measures of wage inequality.

the dynamics of the economy when subjected to aggregate shocks. However, our paper shares the spirit of the papers built on the Lucas-Prescott island model and thus we view our paper as complementary to theirs.

Recent empirical work by Poletaev and Robinson (2008), Gathmann and Schonberg (2010) and Robinson (2010) argue that workers who switch occupations tend to switch to occupations using similar tasks as their previous occupations by constructing measures of task-similarity across occupations. Furthermore, Kambourov and Manovskii (2008) highlights the importance of occupational mobility as they show that there was a high and substantial increase in worker mobility in the United States over the 1968-1997 period. At their most disaggregated measure of occupations they found that the average annual level of occupational mobility was 18% during their sample period. Moscarini and Thomsson (2007) present observations about occupational mobility at the monthly frequency that is consistent with the findings that there have been large swings in the rate of occupational switching. Importantly, they show that there was both a large rise in occupational mobility and a large fall from the mid-1990s until the mid-2000s. At the same time, there was a fall in the unemployment rate as occupational mobility rose and then a rise in the unemployment rate as occupational mobility rates fell. Although, typically unemployment and occupational mobility rates are positively correlated in our model, it does generate extended periods in which the two rates are negatively correlated as seen in the data.

From the point of view of studying occupational mobility and wage inequality, our work is most related to Kambourov and Manovskii (2009a). In addition to presenting some striking empirical observations on the rate of occupation switching its relationship with wage inequality, their paper presents a structural model calibrated to understand the change in wage inequality and occupational mobility between the period 1970-1973 and 1993-1996. Their model presents one calibration which targets the rate of occupational mobility for the 1970-1973 period and another calibration for similar targeted variables over the 1993-1996 period. They then check to see how well their model fits measures of wage inequality for these two periods. An important feature of their model is that they tie the relationship between wage inequality and occupational mobility to a mechanism which generates wage increases due to occupational tenure. Kambourov and Manovskii (2009b) shows that occupational tenure is more important that job tenure as a determinant of wage growth. While extremely similar in spirit, our paper focuses on presenting the dynamics of an aggregate economy constantly being buffeted by aggregate shocks and so we concentrate on inter-occupational wage differences and omit the intra-occupational wage mechanism featured in Kambourov and Manovskii (2009a).

On the front of Lucas-Prescott island models with aggregate shocks, a recent paper by Veracierto (2008) embeds a Lucas-Prescott island model into a real business cycle model with a labour force participation margin and shows that for reasonable parameter values, the model yields procyclical unemployment rate dynamics which are at odds with data from the U.S. economy. While Veracierto's model allows for aggregate productivity shocks, from a modeling standpoint, our model differs in that the choice of occupations by workers is directed whereas Veracierto assumes that workers are randomly assigned to new occupations when they choose to switch occupations. The directed nature of occupation changes in our paper lines up with the recent work on occupation switching and gives rise to persistent unemployment rate and average labour productivity dynamics.

Aside from Veracierto (2008), the only other extension of the Lucas-Prescott island model that we note which allows for aggregate shocks is that by Gouge and King (1997). They present a two-island version of the Lucas-Prescott island model which yields analytical tractability. However, the cost of having only two islands means that there is no concept of distance between islands or occupations. Our paper can then be thought of as complementary to the efforts of Gouge and King (1997) in that we add the concept of distance between occupations in order to address the recent findings concerning individual occupational mobility.

From the point of view of the qualitative properties of the model, Barnichon (2010) uses a New Keynesian model with labour search to address the observation that the correlation between output-per-worker and the unemployment rate has experienced a sign change across the pre-1984 period and the post-1984 period. His model incorporates shocks to total factor productivity and to the money supply which serve as two shocks that have qualitatively different impulse responses on aggregate labour productivity and the unemployment rate. In order to accommodate the change in the correlation between labour productivity and the unemployment rate, he calls for a structural change in the cost of posting vacancies, the responsiveness of monetary policy to technology shocks and the elasticity of labour supply in the post-1984 period. Gali and van Rens (2010) also noted the observation of the changing procyclicality of labour productivity and use a labour search model to understand this observation along with the observations of rising relative volatilities of employment and real wages in the U.S. data. In order to understand these changes through the lens of their search model, they rely on a change in the flexility of labour markets during the mid-1980s. In contrast, our model uses history dependence of the distribution of workers in order to generate the changing relationship between labour productivity and the unemployment rate. As the history of the distribution of workers across occupations depends on the history of shocks as well as the initial distribution of workers, our model does not guarantee that this relationship occurs for any given simulation; however, it consistently happens.

We now lay out a map for the rest of the paper. In Section 1 we briefly present some observations from the U.S. economy that we will subsequently use to check the performance of our model. Next, in Section 2 we set out the baseline model and then in Section 3 we present the main quantitative results of the model. Then a discussion is provided in Section 4 and finally concluding remarks are offered in Section 5.

1 Stylized Facts

In this section we will briefly review three properties of data from the U.S. economy that we will try to understand using our model of occupational mobility. Specifically, we review the observation that there have been qualitative changes in the relationship between average labour productivity and the unemployment rate over the last sixty years. The next observation that we review is referenced from Kambourov and Manovskii (2009a) and relates measures of wage inequality to occupational mobility. Specifically, they showed that between the early 1970s through the early 1990s the U.S. experienced a large increase in both wage inequality and occupational mobility. Lastly, we note that from the early 1980s through to 2006, the unemployment rate and the rate of occupational mobility appeared to be negatively correlated.

Figure 1 provides a plot of the time series for the unemployment rate and average

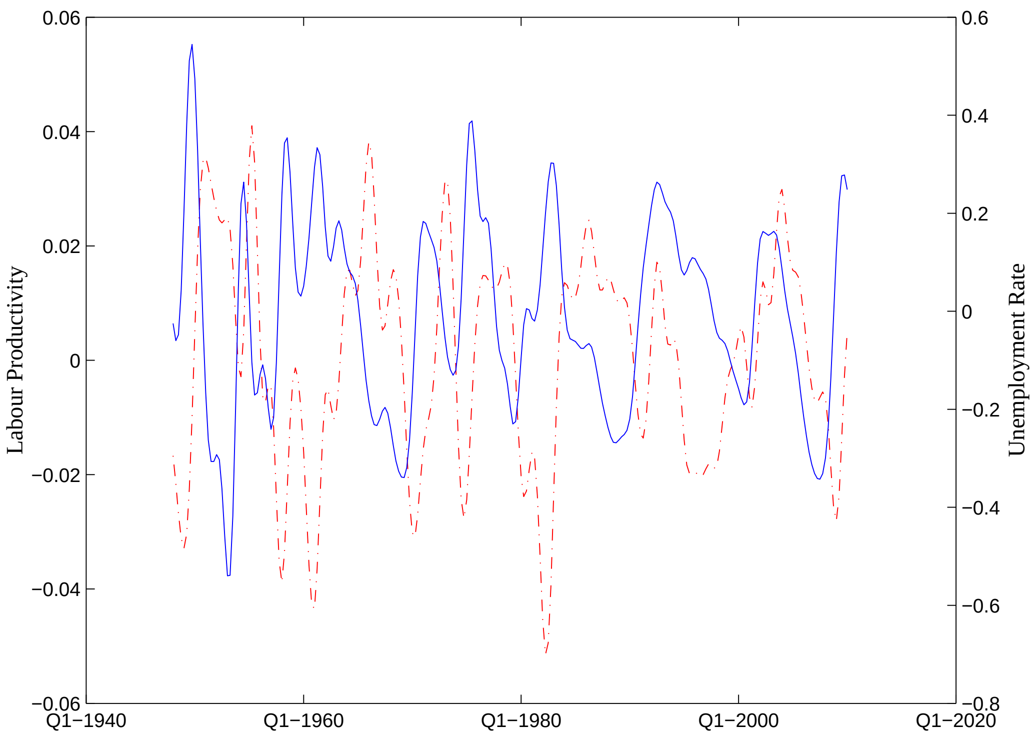

Figure 1: Ave Labour Productivity Vs. Unemployment Rate

Note : The logarithm of the data are taken and detrended using the band pass filter isolating frequencies between 6 and 80 quarters. The broken line plots output per worker and the solid line plots the unemployment rate.

labour productivity after passing them through a band pass filter to extract fluctuations at medium-run frequencies. We focus on medium-run dynamics as we will provide a theory of how occupational mobility affects the relationship between unemployment rates and average labour productivity. Given that most workers who switch occupations likely do so with longer planning horizons than a few years, we find it reasonable to focus on medium-run dynamics.[2] Using ocular econometrics a comparison of the two time series shows that a negative contemporaneous correlation becomes positive sometime in the mid- to late-1980s and arguably the two series move together in the 1970s.

In Table 1 we show the contemporaneous unconditional correlation between average labour productivity and the unemployment rate from the beginning of the 1950s through the end of 2010 broken down by decades.[3] It is clear that there has been a change in the qualitative relationship in the dynamics between the unemployment rate and average labour productivity.[4] There has been some recent discussion that some stylized facts concerning the behaviour of average labour productivity at business cycle frequency are not robust to the measure of labour productivity used.[5] The time series for average labour productivity often used is constructed by seasonally adjusting the ratio of quarterly real non-farm business output (constructed by the Bureau of Labor Statistics (BLS) from NIPA accounts) to employment where the employment series is constructed by the BLS from the Current Employment Statistics (CES). Hagedorn and Manovskii (2010) show that there is a large discrepancy in measures of average labour productivity which are based on employment measures obtained from the CES versus employment measures using employment measures obtained from the Current Population Survey (CPS).

In order to ensure that the qualitative behaviour between average labour productivity and the unemployment rate are not an artifact of using a particular data set, Table 1

[2]In our quantitative exercise, we will use findings by Robinson (2010) and consider occupation changes which require some change in tasks. From a theoretical sense this means that individuals will bear some utility cost from switching occupations.

[3]The results are not qualitatively different if we use 2 quarters as the lower bound of defining our frequencies in lieu of 6 quarters.

[4]See Barnichon (2010) for a more detailed account of this stylized fact.

[5]For example, Hagedorn and Manovskii (2010) show that using CES-based data to construct the average labour productivity series produces different correlation between labor market tightness measures and average labour productivity relative to a series of average labour productivity constructed using CPS-based data.

Table 1: Correlation Between Ave Labour Productivity and the Unemployment Rate

Filtering Full Sample	1950s	1960s	1970s	1980s	1990s	2000s
CES Productivity	-0.4761	-0.5476	-0.2480	-0.6730	0.2902	0.6722
CPS Productivity (All)	-0.6517	-0.9248	0.2675	-0.4716	-0.0102	0.8211
CPS Productivity (NFB)	-	-	-	-0.7201	-0.0023	0.6871

Note : The logarithm of the data is taken and detrended using the band pass filter. For medium run freqencies we isolate freqencies between 6 and 80 quarters. In this table we pass the entire sample through the band pass filters and then using the filtered data we calculate the correlations for each decade. The last decade only uses the filtered data from 2000Q1 to 2008Q1 because endpoint inaccuracies of the band pass filter.

displays the correlation between these two variables over the last six decades using three measures of labour productivity. The first row measures average output per worker where the measure of employment is produced from the CES. This data is obtained from the BLS website. The second row uses average output per worker where the output for all industries is NIPA GDP and employment includes all industries based on CPS measures. Lastly, the third row uses output from the non-farm business sectors measured by the BLS Labor Productivity and Costs program and employment is also for non-farm business sectors as measured from CPS monthly files.[6]

Using CES average labour productivity measures we see that the contemporaneous correltation between average labour productivity and the unemployment rate seems to have switched signs in the later part of the sample, which runs from 1948Q1 through to 2010Q1. Using the CPS measure of labour productivity which includes all sectors, we see that in contrast to using the CES measure, the 1990s exhibits a weakly negative correlation and there is a large qualitative difference between the two measures for the correlation in the 1970s. Nonetheless, the point to note is that there are qualitative sign switches in the correlations over the sample. Finally, if we concentrate on average labour productivity measures using only non-farm business sectors constructed from the CPS monthly files (for which data is available from 1976Q1 onwards) we see again that the strong negative correlation from the 1980s becomes weakly negative in the 1990s and is positive in the

[6]The data sets for CPS Productivity (All) and CPS Productivity (NFB) were obtained from the website of Iourii Manovskii. Details of the data construction are described in Brugemann, Hagedorn, and Manovskii (2010).

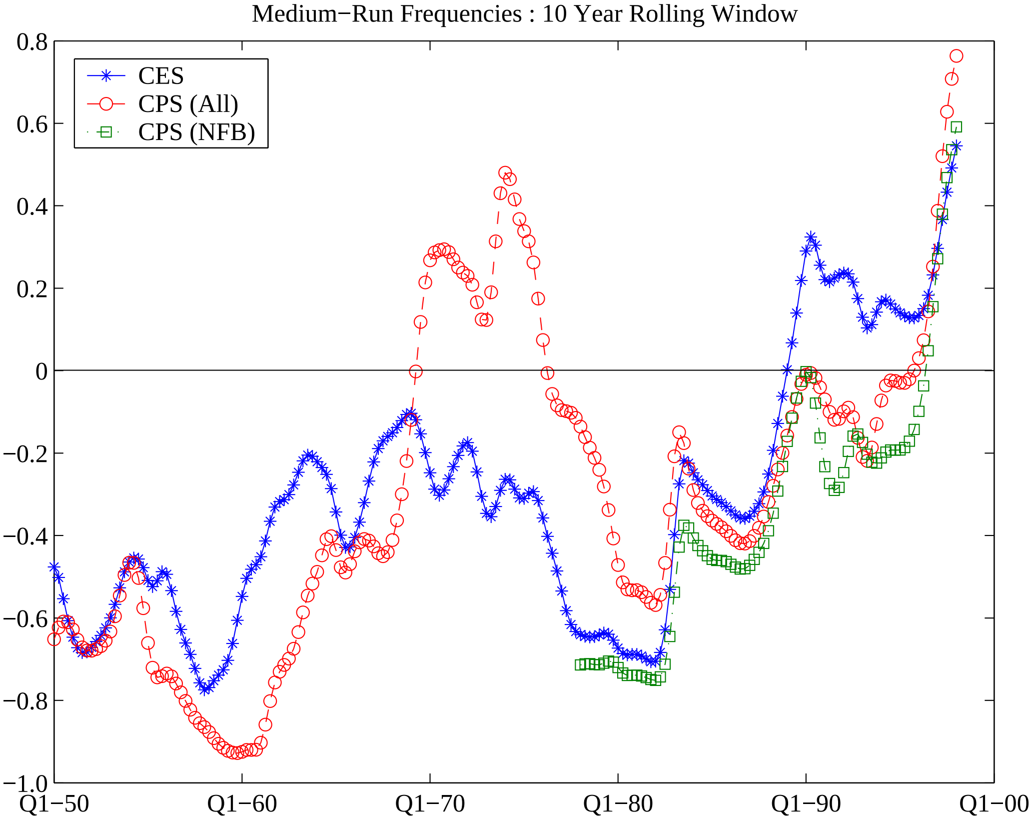

Figure 2: Ave Labour Productivity Vs. Unemployment Rate

Note : The logarithm of the data are taken and detrended using the band pass filter. For medium run frequencies we isolate frequencies between 6 and 80 quarters. In this figure we pass the entire sample through the band pass filters and then using the filtered data we calculate the correlations for each decade. The last decade only uses the filtered data from 2000Q1 to 2008Q1 because endpoint inaccuracies of the band pass filter.

2000s. In order to provide a more comprehensive feel for the data, Figure 2 plots the ten-year rolling window for the unconditional correlation between average labour productivity and the unemployment rate.

While it may appear as though there has been a structural change in the data generating process producing this data, we will argue that these periodic changes in the correlation between the unemployment rate and average labour productivity may be the result of shifts in technology which favour certain occupations over others.[7] The slow moving process of

[7]In an attempt to explain these observations, albeit at business cycle frequencies, using a Mortensen-Pissarides search model, Barnichon (2010) calls for structural changes in monetary policy post-1984 as well

workers retraining to take on new tasks required by switching occupations works together with the direction of technological innovations to result in an evolving statistical relationship between the unemployment rate and average labour productivity of the aggregate economy.

As the dynamics of our model hinge on the back and forth between the incentives that the wage structure across occupations provide individuals to switch occupations and the effects that the distribution of workers across occupations has on the wage structure across occupations, we present some observations concerning wage inequality and occupational mobility. Table 2 reproducess observations from Kambourov and Manovskii (2009a). The first row for each measure of wage inequality shows the authors' measure of wage inequality while the second row controls for years of schooling and experience. The main observation to be taken away from this table is that across three standard measures of wage inequality, there was a substantial increase between the decades of the 1970s and the 1990s. They show that most of this rise occurs in the late-1970s and the early- and mid-1980s. Simultaneously there was an increase in occupational mobility at the 3-digit occupational coding level with the rise in occupational mobility and wage inequality appearing positively correlated over the decade. The same pattern in occupational mobility has been established by Kambourov and Manovskii (2008) using PSID data, Moscarini and Thomsson (2007) using monthly CPS data and Robinson (2010) using annual CPS data. In addition to the rise in occupational mobility, Moscarini and Thomsson (2007) also show that there was an equal size decline in occupational mobility after the mid-1990s. This suggests that there may be important medium-run dynamics in the rate of occupational switching and that these changes are not necessarily results structural change but rather the slow response of the economy to shocks which is the argument we will push in our paper. Of course, our story may be just one of many reasons behind such shifts and we do not take a stance on ruling out other explanations in this paper.

We will try to show that our model can generate switching in the correlation between the unemployment rate and average labour productivity over time. Also we will argue that our

as to household preferences and the cost of job creation faced by firms. Gali and van Rens (2010) argue that changes in labour market flexibility over the period may help understand the changes in these correlations. We do not argue that there were no changes in the structure of the economy over this sample which may affect the observed correlations. Rather we provide a different mechanism to help think about the observed properties of the data.

Table 2: Wage Inequality and Occupational Mobility

		1970-1973	1993-1996	% Change
Variance of Log Wages	Overall	0.225	0.354	57%
	Within-Group	0.162	0.248	53%
Log 90/10 Ratio	Overall	1.167	1.448	24%
	Within-Group	0.975	1.192	22%
Gini Coefficient	Overall	0.258	0.346	34%
	Within-Group	0.215	0.273	27%
Occupational Mobility		0.157	0.205	31%

Notes: Statistics taken from Tables 1 and 3 of Kambourov and Manovskii (2009a). Details on construction of these statistics can be found in Kambourov and Manovskii (2009a) and Kambourov and Manovskii (2008).

model is able to generate strong positive correlations between measures of wage inequality and the rate of occupation switching. In order to tie these two properties together we will see whether the model can account for the observation that the unemployment rate and the rate of occupational mobility was negatively correlated over the period from 1979 to 2006. Figure 3 plots monthly observations on the rate of occupational mobility as calculated by Moscarini and Thomsson (2007) against the unemployment rate. Due to changes in the methods used to construct the measure of occupational mobility, the data from the subsamples 1979-1982, 1983-1993 and 1994-2006 are not directly comparable. Nonetheless, we get a notion of the relationship between the two variables by examining this data. If we calculate the correlation of the two series for the periods we obtain $(-0.237, -0.432, -0.134)$ respectively. Broadly we can say that these correlations are negative.[8] Within the subsample from 1994-2006 the data is constructed in a consistent manner. From February 1994 to November 2002 the correlation between the unemployment and occupational mobility was 0.212 while between January 2003 and April 2006 the correlation was -0.507. While this is a small sample of data and we acknowledge that care need be taken in generalizing from these observations, we will see if the model can generate such correlations in order to put more discipline on the way that we view the implications of the model.

Summarizing, after looking at the data, we will be interested to see if our model can

[8] We do not pass this data through the band pass filter to isolate medium frequency movements because due to endpoint inaccuracies of the band pass filter we would need to drop the data from the first two years and last two years of each subsample leaving very little data to calculate the correlations.

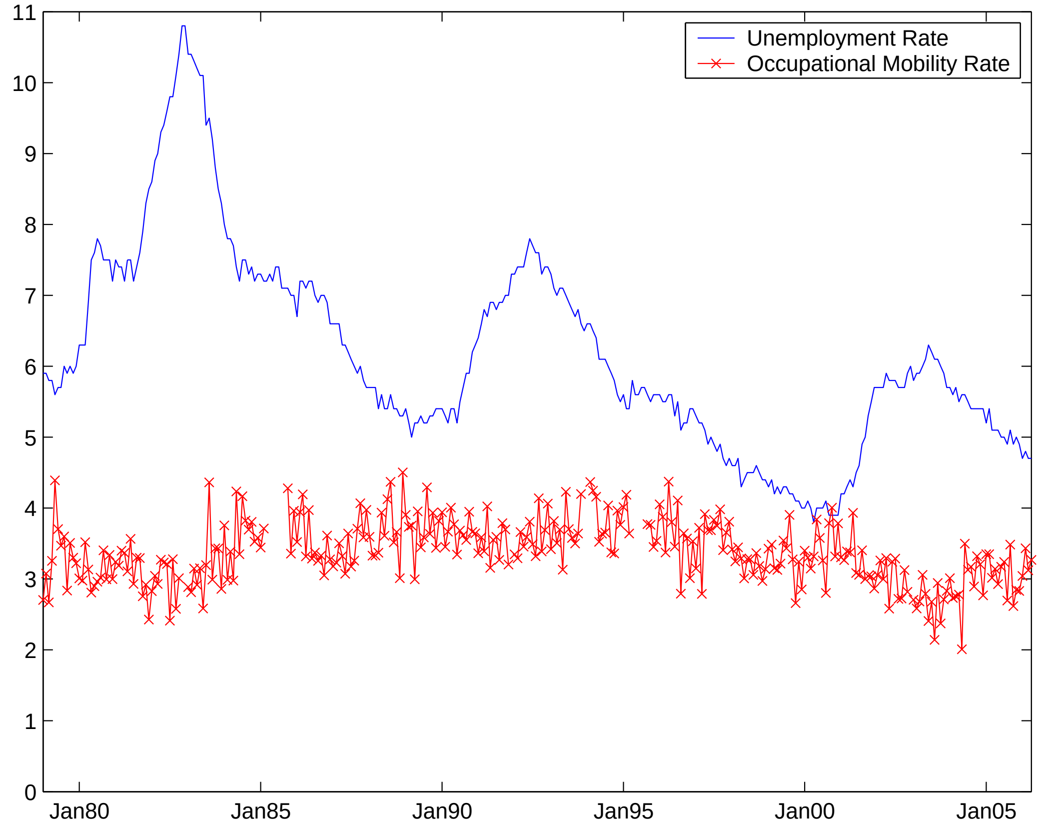

Figure 3: U.S. Economy : Unemployment Rate and Occupational Mobility

produce simulated data such that (i) the correlation between the unemployment rate and average labour productivity can change signs across decades within a simulation, (ii) there can be decades in which the correlation between the unemployment rate and the occupational mobility rate is negative, and (iii) there can be periods of 20 years during which there are large differences between the maximum and minimum levels in measures of wage inequality and the occupational mobility rate with the extra constraint that wage inequality and occupational mobility are positively correlated.

2 The Model

We construct a model in which technological innovations constantly shifts relative labour productivities across occupations. Workers must choose whether to incur the costs of retraining to work in new occupations, being unemployed if no work is available in his/her current occupation or working at the prevailing wage his/her current occupation. As a simplification, we assume workers only possess the skills to work in one occupation at any point in time. The timing is as follows: at the beginning of the period, workers are distributed across a continuum of occupations. Each worker draws an idiosyncratic fixed cost of switching occupations in the current period. The distribution of labour productivity across occupations in the current period is then revealed. At this point workers decide whether to switch occupations, or stay in their current occupation. After these decisions are made, production, trade and consumption occurs. Finally, the current period ends. In the following subsections we provide more explicit details regarding the economy's production structure, the problem of the individual worker and the definition of an equilibrium.

2.1 Technology

There is a continuum of occupations on a circle with radius one indexed by $i \in [0, 2\pi]$.[9] Consider a uni-modal function $g(\cdot)$ that is symmetric around the mode with domain $[-\pi, \pi]$. Let θ denote the location of the mode on the circumference in the current period and assume

[9] In lieu of occupations, locations on the circumference of the circle may be thought of as geographic locations or, more abstractly, geographic-specific occupations.

that θ follows the process[10]

$$\theta' = \begin{cases} \theta + \epsilon' & \text{if } \theta + \epsilon' \in [0, 2\pi] \\ \theta + \epsilon' - 2\pi & \text{if } \theta + \epsilon' > 2\pi \\ \theta + \epsilon' + 2\pi & \text{if } \theta + \epsilon' < 0 \end{cases}, \qquad \epsilon \sim F[-\pi, \pi].$$

Let $F(\epsilon)$ be continuously differentiable on $(-\pi, \pi)$ and denote its density by $f(\epsilon)$. Let the relative location of $i \in [0, 2\pi]$ from the location of the mode at time t be given by

$$\delta(i) = \min\{\theta - i, 2\pi - \theta + i\}$$

such that the distance of i from θ is $|\delta(i)| = \min\{|\theta - i|, |2\pi - |\theta - i||\}$. Let the value $A(i) = g(\delta(i))$ be the level of labour productivity in occupation i given that occupation i has a location $\delta(i)$ relative to the location of the productivity frontier. We assume that $g(\cdot)$ is twice continuously differentiable. Denote the height of $g(\cdot)$ at the mode by $\overline{A}$ so by construction, $g(0) = \overline{A}$. Finally, assume that $g(\cdot) \in [0, \overline{A}]$ with $\lim_{\delta \to -\pi} g(\delta) = \lim_{\delta \to \pi} g(\delta) \cong 0$.

Figure 4 shows an example in which the mode shifts from one period to the next while the function $g(\delta)$ retains its shape. Figure 5 depicts an example of the determination of $\delta(i)$ for occupation i when θ shifts over time.

2.2 Production

2.2.1 Final Good Firms

There is an aggregate consumption good which is produced by perfectly competitive final good firms. The final good firms use the output of intermediate good firms as inputs. Intermediate good i can only be produced by firms in occupation i where $i \in [0, 2\pi]$. Final good firms take the price of intermediate goods from occupation i as given and the price of the final good is normalized to one. Notice that given the assumptions on productivity, the output of a firm-worker pair is independent of the occupation's identity; it is only dependent on the distance of the occupation from the most productive occupation. Thus, for an occupation that is δ from the most productive occupation, we can write the output of a worker-firm pair as $A(\delta)$. Denote the measure of workers in occupation δ in the current period by $\psi(\delta)$. If there is production from the occupation δ from the frontier in the current period, then total output from this occupation, is denoted by $y(\delta, \psi) = A(\delta)\psi(\delta)$.

[10] We follow the convention that "primed" variables indicate the value of the variable in the next period.

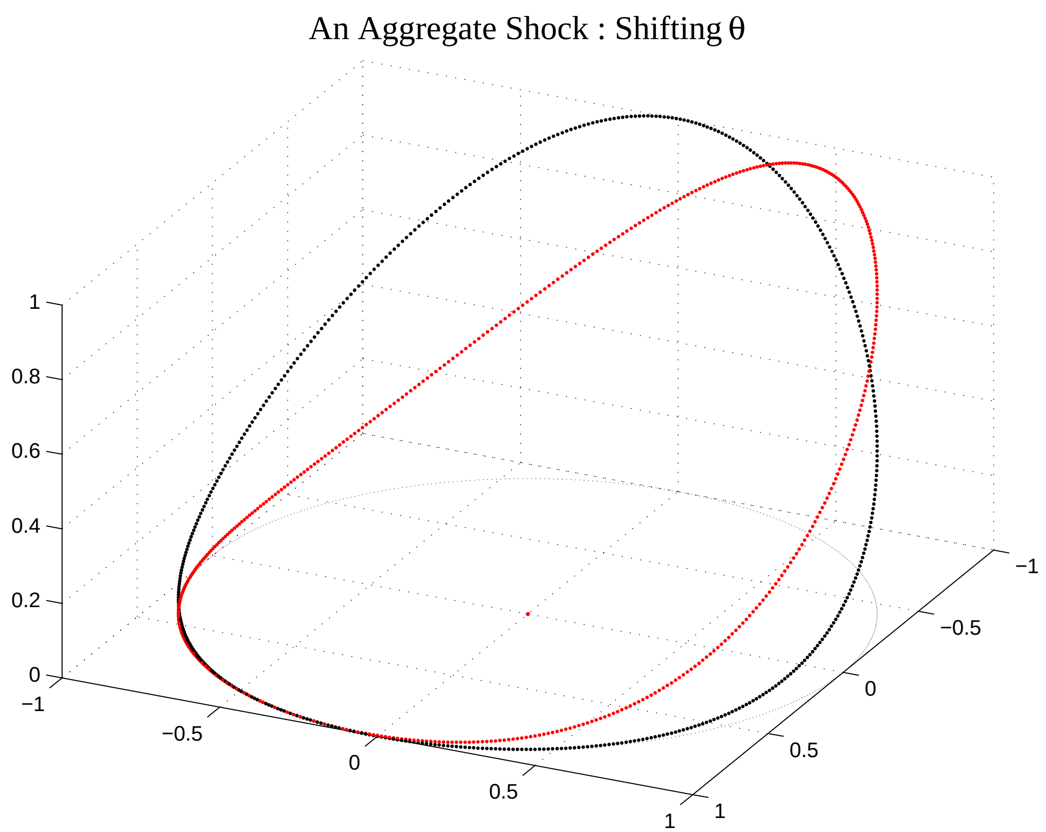

Figure 4: An aggregate shock : shifting the mode of $g(\delta)$

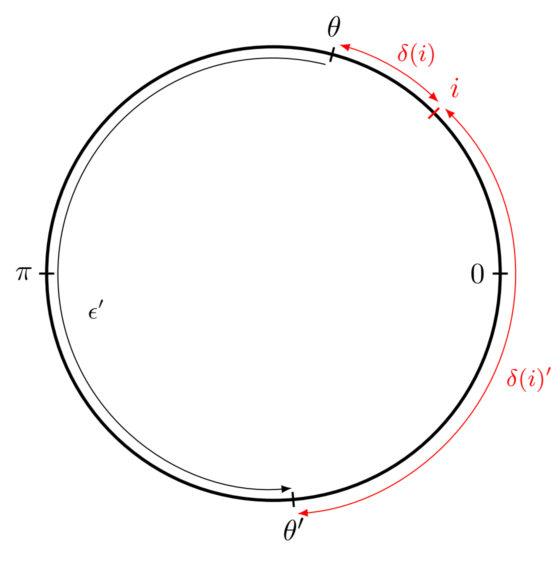

Figure 5: The determination of $\delta(i)$

Final good firms aggregate across all available intermediate goods using the production function

$$y(\psi) = \left[\int_{\Omega} y(\delta, \psi)^{\frac{\chi-1}{\chi}} d\delta\right]^{\frac{\chi}{\chi-1}}$$

where Ω is the set of all occupations with strictly positive output in the current period, and $\chi > 0$ is the elasticity of substitution between intermediate goods.

2.2.2 Intermediate Good Firms

Intermediate good i is produced by worker-firm pairs situated at location $i \in [0, 2\pi]$ through use of a linear production technology, such that, in the current period, each employed worker in occupation i produces an amount $A(i)$. Again, given the modeling assumptions we can rewrite this in terms of intermediate goods δ from the frontier so that $\delta \in [-\pi, \pi]$ and output of a worker-firm pair in an occupation δ from the frontier in the current period is $A(\delta)$. There are no set-up costs of period production which means that as long as production is profitable, firms continue to enter any given occupation until all workers in that occupation are hired. Suppose that workers have an outside option relative to employment of b if they stay at home. This period payoff to unemployment, b, can be either be thought of as output from home production or as unemployment benefits. Thus the minimum period operating costs to a firm is b as firms have to cover the participation wage of their workers.

Let the price of intermediate good produced by the occupation δ from the frontier, given the distribution of workers $\psi(\delta)$ be $p(\delta, \psi)$. All intermediate good firms take the price of their output as given. The inverse demand function for any firm operating in an occupation δ from the frontier is given by

$$p(\delta, \psi) = \left[\frac{y(\delta, \psi)}{y(\psi)}\right]^{\frac{-1}{\chi}} \qquad \forall \delta \in \Omega,$$

where we have normalized the price of the final good to be one.

2.3 Worker's Problem

Now we consider the problem of a worker. There are no savings in this economy and a worker's period utility from consumption is given by $u(c)$ with $u' > 0$, $u'' < 0$ and

$u(0) > -\infty$. A worker who is located in occupation i may choose to be employed, if there are firms is occupation i, unemployed but attached to occupation i, or not participating in the labour force. Period consumption in the unemployed state is $b \geq 0$. When a worker is not in the labour force, the worker may choose to move to a different occupation. When choosing to switch occupations a worker enjoys consumption in the amount b but must incur an idiosyncratic fixed utility cost of switching, z. Each period a worker's idiosyncratic cost of moving is drawn from a time invariant distribution $H(z)$ with density $h(z)$ and support $[0, \infty)$. Assume that $H(z)$ is twice continuously differentiable everywhere and that $\int_0^\infty z dH(z)$ is bounded from above. Furthermore, there is also convex disutility cost of retraining, $\varphi(x) \geq 0$, where x is the measure of distance between the worker's current occupation and his/her desired, new occupation. Simply, x is the distance around the circumference of the occupation ring between the old and the new occupation. Assume that $\varphi(0) = 0$ so the cost of moving zero distance is zero. Each worker faces the probability of death, $1 - \beta$, at the each period. In order to keep the size of the working population at one, a measure $1 - \beta$ of new workers are born each period. New born workers are uniformly distributed across all occupations.

When a worker is employed in occupation δ from the frontier, the period wage is $w(\delta, \psi)$ which yields period utility $u(w(\delta, \psi))$. Let the value of being employed in a occupation that is δ units away from the productivity frontiers be $V(\delta, \psi)$. Denote the value of being unemployed in a occupation δ from the frontier to be $U(\delta, \psi)$ and denote the value of being a non-participant or occupation switcher (and retraining) to be $T(\delta, \psi)$. In this model there is no "search unemployment" as all workers who are unemployed are not technically looking for jobs because they know that there are no jobs available in their occupation. In some sense, all workers who are unemployed have been displaced and are not looking for a job so in the spirit of Alvarez and Shimer (2009a), we refer to these workers as "rest unemployed". Due to the nature of unemployment in our model, our paper can be viewed as complementary to papers in which unemployment arises due to search frictions.[11] We

[11] For example, papers such as Burdett, Shi, and Wright (2001), Peters (2000) and Shimer (2005) provide a theory of unemployment due to coordination frictions between workers and firms with job vacancies. Models of competitive search as in Moen (1997) and the popular random matching models described in Pissarides (2001) are less specific regarding the nature of the matching frictions in the labour market and may capture the structural unemployment that we model explicitly in our paper.

refer to workers in the retraining state as non-participants because they are not actively seeking a job.

It is clear that as competition amongst firms results in firms paying at least the participation wage, workers will never choose to quit a job to enter unemployment; if workers quit then they must be choosing to retrain. Therefore the only time a worker enters the unemployment pool is if the worker is displaced, that is, firms in the given occupation shut down and the worker does not choose to retrain. However, workers may choose to quit in order to retrain.

Suppose that $\Omega := \{\delta : y(\delta) > 0\}$, so that Ω is the set of all occupations in which production is strictly positive. Then we can write

$$\begin{aligned} V(\delta,\psi) &= u(w(\delta,\psi)) + \beta \left\{ \int_{[-\pi,\pi]\setminus\Omega'} \int_0^\infty \max\left\{T(\delta',\psi') - z', U(\delta',\psi')\right\} dH(z')dF(\delta'|\delta) \right. \\ &\quad \left. + \int_{\Omega'} \int_0^\infty \max\left\{V(\delta',\psi'), T(\delta',\psi') - z', U(\delta',\psi')\right\} dH(z')dF(\delta'|\delta) \right\}. \qquad (1) \end{aligned}$$

Workers who are unemployed are faced with the decision to stay unemployed in their current occupation or to exit the labour force and retrain. The value function for an unemployed worker in a occupation δ away from the frontier is

$$\begin{aligned} U(\delta,\psi) &= u(b) + \beta \left\{ \int_{[-\pi,\pi]\setminus\Omega'} \int_0^\infty \max\left\{T(\delta',\psi') - z', U(\delta',\psi')\right\} dH(z')dF(\delta'|\delta) \right. \\ &\quad \left. + \int_{\Omega'} \int_0^\infty \max\left\{V(\delta',\psi'), T(\delta',\psi') - z', U(\delta',\psi')\right\} dH(z')dF(\delta'|\delta) \right\}. \qquad (2) \end{aligned}$$

Finally, workers who are non-participants will incur a fixed cost to retrain. The retraining decision comes at a utility cost of z in addition to a convex utility cost, $\varphi(x)$, in the distance traveled, x, which denotes the measure of occupations that the worker passes over in the period. Thus, for non-participants, $\delta' = \delta + x + \epsilon$ allowing the value function for a non-participant who is currently δ from the productivity frontier to be written as

$$
\begin{aligned}
T(\delta,\psi) &= \max_x \{u(b)-\varphi(x) \\
&\quad +\beta\left\{\int_{[-\pi,\pi]\setminus\Omega'}\int_0^\infty \max\left\{T(\delta',\psi')-z',U(\delta',\psi')\right\}dH(z')dF(\delta'|\delta+x)\right. \\
&\quad \left.+\int_{\Omega'}\int_0^\infty \max\left\{V(\delta',\psi'),T(\delta',\psi')-z',U(\delta',\psi')\right\}dH(z')dF(\delta'|\delta+x)\right\}.
\end{aligned}
\tag{3}
$$

Note that the $T(\cdot)$ is the value of switching occupations after having incurred the fixed cost of switching occupations.

2.4 Equilibrium

Definition 1 *An equilibrium is*

1. *a set of value functions $V(\delta,\psi)$, $U(\delta,\psi)$, and $T(\delta,\psi)$, along with the decision rule for relocation, $x(\delta)$, and choice of employment states that satisfy the Bellman equations (1), (2), and (3),*
2. *a distribution function $\psi(\delta)$ consistent with the individual's decision rules and satisfying a transition function $\psi' = \Gamma(\psi,\epsilon')$,*
3. *a price function $p(\delta,\psi)$ that clears all markets in which there is output, and*
4. *where there is production, a wage function $w(\delta,\psi)$ such that firms always obtain zero profits.*

Free entry into each occupation causes the worker's wage to be equal to total revenues of the firm, $p(\delta,\psi)A(\delta)$, so that $w(\delta,\psi) = p(\delta,\psi)A(\delta)$. Free entry implies that firms in occupation $\delta \in \Omega$ pay their workers a wage of

$$
w(\delta,\psi) = A(\delta)^{\frac{\chi-1}{\chi}}\psi(\delta)^{-\frac{1}{\chi}}y(\psi)^{\frac{1}{\chi}}.
$$

Notice that as we push χ to infinity, the wage becomes equal to labour productivity only; the measure of workers in an occupation no longer matters for the wage. We will refer to this case as the partial equilibrium model.

3 Quantitative Results

In this section we will first describe the choice of parameter values followed by a presentation of the quantitative behaviour of the model conditional on the chosen parameter values. This is followed by an exercise in which we contrast some impulse responses to reallocation shocks in order to clarify the main mechanism at work in the model economy. Then we provide a partial equilibrium model in order to highlight certain properties of the general equilibrium model.

3.1 Parameterization

The model is highly stylized and parsimonious in the number of parameters to be calibrated for the following reason. Several features of the model do not have counterparts in much of the previous literature so there are no existing values for some parameters for which we can easily adopt from other work. Also, as the distribution of workers across occupations affects wages which feeds back into mobility decisions, the solution procedure is computationally intensive. In particular, as is explained below, the steady state of the model economy is, in some sense, a "near best" state for the unemployment rate. This renders the usual procedure of calibrating parameters by targeting steady state statistics an undesirable procedure and results in a large increase in the time required to choose hit targeted statistics by choice of a set of model parameters.[12]

In choosing the model's parameter values for the following computational exercise, we stray from the standard method of matching steady state targets. The reason for doing so is that the steady state is close to the lower bound of feasible unemployment rates and occupational switching rates in this economy. Intuitively, if the economy starts off in a stochastic steady state (i.e. a distribution of workers across occupations when the economy is continually hit with zero shocks) then any reallocation shock will result in an increase in both the unemployment rate and the occupational switching rate. Given this obstacle, we choose parameter values such that specific targets lie in the 95% band from an ensemble of simulations.[13] As it takes between several hours and several days to solve the model

[12] The computational procedure used to approximate the model is provided in the Appendix.

[13] For the current version, for a given set of parameters we solve the nonlinear model and simulate the economy 1000 times with 249 periods per simulation after an appropriately long burn-in period is discarded.

Table 3: Parameter Values

Parameter	Value	Description
β	0.99375	Subjective discount factor
σ	1	Elasticity of intertemporal substitution
χ	3	Elasticity of substitution in final goods CES aggregator
$\underline{a}$	0.01	Lower bound for labour productivity
$\overline{a}$	1	Upper bound for labour productivity
ξ	2.25	Shape parameter for labour productivity function $g(\delta)$
α_ϵ	120	Parameter governing distribution of shocks
μ_z	4.75	Mean of idiosyncratic shocks from Gamma Distribution
σ_z^2	2.75	Variance of idiosyncratic shocks from Gamma Distribution
b	2.70	Fraction of $\overline{a}$ enjoyed by unemployed
η	60	Weight on quadratic retraining cost

for a given set of parameters we loosely target particular statistics observed in U.S. data described below.

Table 3 displays the parameter values used in the numerical simulations. In the simulations that follow, a period represents three months so that the value of β is such that a worker is expected to have a work life of forty years. We choose utility from consumption to be given by $u(c) = \frac{c^{1-\sigma}-1}{1-\sigma}$ and fix $\sigma = 1$ so that utility is logarithmic. The maximum and minimum values of labour productivity are normalized to be $\overline{a} = 1$ and $\underline{a} = 0.1$. As long as $b > \underline{a}$ the lower bound of labour productivity plays no role in the analysis. As the final goods sector aggregates across intermediate goods produced by different occupations, there is little guidance on the choice of χ. Alvarez and Shimer (2009a) argue that in a model of sectoral mobility elasticities in the range of 2.8 to 4.5 may be reasonable. Our model is not directly comparable to theirs as we consider occupational mobility and they consider sectoral mobility. Kambourov and Manovskii (2009a) consider a model in which all occupations produce a homogeneous good so the analogy in our model would be that intermediate goods would be perfectly substitutable. This would reduce our model to a partial equilibrium model in which there would be no feedback from the distribution of workers to the distribution of wages. In an attempt to obtain some comparability between

Then, for a particular statistic, we sort the statistics in our ensemble of 1000 data points and chop off the lowest and highest 2.5%. We then check to see if the statistic from the U.S. data lies in this interval from our simulated ensemble.

our work and other work using the Lucas-Island model set-up we choose to present results for $\chi = 3$ and $\chi = 2.75$.

For simplicity, the shape of the function, $g(\delta)$, is chosen to be uni-model, symmetric around zero and differentiable almost everywhere. We employ the beta distribution which typically is governed by two parameters. However, for the beta density to be symmetric the two parameters must be equal which reduces the number of parameters to be calibrated. Denoting the shape parameter for the $g(\delta)$ function to be ξ we choose ξ to be equal to 2.25. The distribution for idiosyncratic fixed costs of retraining is chosen to be gamma with mean and variance denoted by μ_z and σ_z^2. Finally, the distribution from which period shocks to θ is also chosen to be a symmetric beta distribution with parameter α_ϵ.

In choosing the values of the parameters, one target that we attempt to approximately match is an average unemployment rate of 5.7% which is the sample average of the U.S. unemployment rate between the first quarter of 1948 and the first quarter of 2010. In recent work, Kambourov and Manovskii (2009a) find that between the early 1970s and the mid-1990s the average annual rate of occupational switching is 18%.[14] They show that in the early 1970s the occupational switching rate was about 16% and rose to approximately 20% by the mid-1990s. In a similar effort, using Current Population Survey data from 1979-2006, Moscarini and Thomsson (2007) argue that at a monthly frequency, the occupational switching across three-digit occupational classification in the US is approximately 3.5%.[15] With this number, if individuals face an independent switching probability from month-to-month, then annual mobility would be 34%. While Moscarini and Thomsson find results consistent with the upsurge in occupational mobility between 1979-1995 found by Kambourov and Manovskii, it is shown that there was a large decline in occupational mobility between 1996-2004 only to see a cyclical rebound after 2004.

[14] Kambourov and Manovskii (2009a) use 3-digit occupation classification codes employed in the U.S. Census and the PSID dataset to measure occupational mobility. A worker is defined as having switched occupations if he reports that he currently works in a different occupation than his last reported occupation in the previous years survey. The sample is restricted to male heads of households, aged 23-61, who are not self- or dual-employed and who are not working for the government.

[15] Moscarini and Thomsson (2007) use the monthly CPS data from 1979-2006 and calculate occupational mobility rates for all male individuals between the ages 16-64 who were employed in consecutive months. Due to changes in the classification system of three-digit occupation codes in 1983 and due to changes in the CPS interviewing methodology in 1994, their sample is broken down into three subsamples (1979-1982, 1983-1993, and 1994-2006) within which comparisons are consistent but across which there are issues of consistency of measurement.

We are interested in choosing parameters to fit median distances between occupation switching as well as frequency of occupation switches and so we prefer to use targets on distance between occupations and rates of occupation switching that are derived from the same data set. Recent work by Poletaev and Robinson (2008) uses information from the Dictionary of Occupational Titles (DOT) to construct distances between occupations. The Dictionary of Occupational Titles contains information that allows them to distinguish between occupations that are similar and those that are different in the underlying tasks they perform. Poletaev and Robinson (2008) use this information on job characteristics and construct low-dimensional skill vectors for each occupation in the Dictionary of Occupational Titles. They then produce measures of distance measures on these skill vectors to interpret distance between any two occupations. Robinson (2010) combines these measures of distance with survey data recorded in the March Current Population Surveys in order to provide a description of the patterns of occupational mobility and distance between occupations involved in switches.

Importantly, Poletaev and Robinson (2008) and Robinson (2010) distinguish between occupation switches which are determined by changes in the occupational digits as classified in the DOT and occupation switches which involve a change in the main tasks (or significant changes in the vector of skills). As they note, some occupation switches as determined by switches in the 3-digit occupational code involve negligible change in the skills involved in the job. In our model, workers who switch occupations are incurring increasing costs in the distance between their current occupation and their new occupation, which is to characterize the costs of learning new tasks inherent in occupation switching. Thus we exploit some of the empirical findings from Robinson (2010) in our baseline parameterization which document the rate of occupational switching when accounting for non-trivial changes in the skill vector associated with two different occupations.

In Robinson (2010) it is found that using 3-digit occupation code changes as the definition of occupation switches, from 1982 to 2001, the rate of occupational mobility was 9.08% in a pooled sample with a minimum annual rate of 7.06% in 2001 and a maximum annual rate of 10.85 in 1988. It is noted that while the level of measured mobility is lower than those found by Kambourov and Manovskii (2009a) and Moscarini and Thomsson (2007) the pattern of occupational mobility is similar in that there is an upward trend in occupational mobility from the early 1980s to the early 1990s and then a downward trend from the early 1990s until the early 2000s. When occupation switches are classified as requiring a change in the main skill in the skill vector with the highest score then the pooled rate of occupational mobility between 1982-2001 drops to 5.71%. In further refinements of these skill-portfolio measures of occupational mobility that use information on the level and distance of changes in the skill portfolio to classify switchers, the rate of occupational mobility falls as low as 3.58%. We use the numbers to give us a rough target for annual occupational switching rates in our simulations in order to help pin down an interesting region of the parameter space.

We use a quadratic function to represent retraining costs to the worker $\varphi(x) = \frac{\eta}{2}x^2$. The parameters are chosen to loosely fit the median distance between occupational switches as reported by Robinson (2010). Recent work by Poletaev and Robinson (2008) uses information from the Dictionary of Occupational Titles to construct distances between occupations. Using these distance metrics, Robinson (2010) argues that using the 3-digit occupation classifications in the Dictionary of Occupational Titles that the median distance in changes for occupational switchers is approximately 0.907 and the maximum distance between occupations (for a 4-dimensional skill vector measure of distance) used by Robinson (2010) is 8.072. Using data from the Displaced Workers Survey, it is found that the median distance of occupational switchers is 0.871. Relative to our model, where the maximum distance between two occupations is π, we target a median distance of approximately 0.11. In our simulations we calculate the median distance in occupational switches for all movers from the first quarter of a year to the first quarter of the next year. Doing so with the parameter values displayed in Table 3 yields an average distance across occupational switchers over 50000 periods of 0.103.

In order to help give some context to the level of consumption enjoyed during unemployment, b is chosen so that on average b is between 33% to 39% times the average wage. This is so that b is in the lower and upper bounds of the replacement ratios paid to unemployed workers in the U.S. Our particular parameterization returns an average replacement ratio of approximately 34% which is towards the lower end of the range for replacement ratios

Table 4: Medium Run Frequencies, Quarterly U.S. Data : 1948Q1 - 2010Q1

	u	Y/e
Standard Deviation	0.2016	0.0191
Quarterly Autocorrelation	0.9520	0.9492
Correlation Matrix		
Unemployment Rate	1.0000	-0.3713
Ave Labour Productivity	-	1.0000

Notes: The logarithm of the data are taken and detrended using the band pass filter isolating frequencies between 6 and 80 quarters. The mean of the relevant statistics are presented in the table with their standard deviations in parenthesis.

offered in the U.S.

Finally, we choose the parameters so as to approximately fit the volatility and persistence of average labour productivity as measured by output per worker in the U.S. data. Here persistence is measured as the correlation between average labour productivity and its first lag.

3.2 The Baseline Model

Table 5 presents statistics generated by the model economy which can be compared to some medium-run frequency data collected for the U.S. economy. Given the parameterization, we can see that all but one of the statistics in Table 4 fit within the 95% intervals of the simulated ensemble. Statistics from Table 4 that were not targeted are the volatility and first order autocorrelation for the unemployment rate as well as the contemporaneous unconditional correlation between the unemployment rate and average labour productivity. An issue to note as highlighted by the ensemble of statistics is that the 95% regions are all large indicating that for this parameterization, the behaviour of the aggregate statistics exhibit quite large fluctuations across samples of 249 periods. Furthermore, the volatility of the unemployment rate seems high relative to that in the U.S. data, which may not be surpising as the model is overly simplistic and omits features of the aggregate economy such as savings which may act to smooth the behaviour of the aggregate unemployment rate, even at medium-run frequencies.

In order to be point out the large nature of the fluctuations, the top panel of Figure 6

Table 5: Simulated Statistics : Medium-Run Frequency

	u	Y/e	n
Standard Deviation	0.3617	0.0176	0.2507
	(0.2492,0.4657)	(0.0121,0.0243)	(0.1680,0.3681)
Quarterly Autocorrelation	0.9377	0.9395	0.9308
	(0.9015,0.9639)	(0.8996,0.9687)	(0.8867,0.9619)
Correlation Matrix			
u	1.0000	-0.6711	0.5936
		(-0.9044,-0.1992)	(0.1853,0.8439)
$\frac{Y}{e}$	-	1.0000	-0.8018
			(-0.9410,-0.5772)
n	-	-	1.0000

Notes: To construct the data the model is simulated 1000 times. The logarithm of the simulated data are taken and detrended using the band pass filter isolating frequencies between 6 and 80 quarters. The mean of the relevant statistics are presented in the table with their standard deviations in parenthesis.

Table 6: Simulated Statistics : Wage Dispersion Measures

	$\sigma^2(\log(w))$	$\log\left(\frac{w_{90}}{w_{10}}\right)$	Gini
Mean	0.4406	0.7470	0.1428
	(0.1872,0.7062)	(0.5692,0.8904)	(0.1158,0.1634)
Standard Deviation	0.4905	0.1804	0.1179
	(0.3407,0.6328)	(0.1244,0.2491)	(0.0731,0.1749)
Quarterly Autocorrelation	0.9432	0.9289	0.9384
	(0.9068,0.9674)	(0.8920,0.9590)	(0.8982,0.9670)
Correlation Matrix	$\sigma^2(\log(w))$	$\log\left(\frac{w_{90}}{w_{10}}\right)$	Gini
u	0.9805	0.6587	0.6704
	(0.9562,0.9927)	(0.2981,0.8734)	(0.2474,0.8917)
$\frac{Y}{e}$	-0.7501	-0.8639	-0.9625
	(-0.9539,-0.3020)	(-0.9560,-0.6927)	(-0.9877,-0.9082)
n	0.7042	0.7474	0.7752
	(0.3433,0.9092)	(0.4941,0.9078)	(0.5385,0.9299)

Notes: To construct the data the model is simulated 1000 times. The logarithm of the simulated data are taken and detrended using the band pass filter isolating frequencies between 6 and 80 quarters. The mean of the relevant statistics are presented in the table with their standard deviations in parenthesis. The exception to this are the statistics for the mean for which we report the mean of the unfiltered data (without taking the logarithm) generated by the model economy.

provides a scatterplot of the maximum and minimum annual rate of occupational mobility for each simulated run in the ensemble of 1000 simulations used to calculate the statistics in Tables 5 and 6. While the mean of the average occupation switching rate across the ensemble of simulations is 5%, most of the simulated sequences of time series have a minimum rate around 2.0% to 3.0% and a maximum occupation switching rate of approximately 8% to 12%. From Robinson (2010) the comparable magnitudes in the United States between 1982 and 2001 were 4.3% for the minimum and 6.8% for the maximum. Thus the baseline calibration generates excessive volatility. A similar exercise for the unemployment rate shows that the minimum unemployment rate across simulations seems to be around 1.25% and 1.75% with a maximum of around 15% to 20%. For comparison, the minumum quarterly averaged unemployment rate for the U.S. in the sample is 2.6% and the maximum is 10.7%. One reason for this is that there is too much substitutability across goods produced by occupations. Similar results have been obtained for a calibrated model with $\chi = 2.75$.[16]

Figure 7 provides a histogram of the unemployment rate from the simulation in the ensemble that best fits the U.S. experience. In selecting the particular histogram of best-fit, we constructed the histogram for each simulation in the ensemble and calculated the mean squared deviation relative to the histogram for the U.S. data. The top two panels display the histogram and the associated cdf for the simulation run from the ensemble with the lowest mean squared deviation. In this particular sample run, the model economy produces more periods of unemployment rates below the minimum observed in the U.S. sample data and a few periods in which the unemployment rate exceeds 10%. The main reason for this is that workers take large changes in occupational distance when they retrain and combined with the frequency of occupational switching with relatively small reallocational shocks, workers bunch up in high paying occupations quickly. When the dispersion of workers across occupations gets low, small shocks can then result in high unemployment rates and high rates of occupation switching. Surprisingly, at this level of complementarity between outputs across occupations, there is still quite a bit of volatility. Lowering the elasticity of substitution acts to lower the volatility in output but also lowers the volatility in labour productivity generating a trade-off when calibrating the model.

[16]Some of the results for the case of $\chi = 2.75$ are shown but as the results are very similar to those of the $\chi = 3$ case, we suppress some of the tables.

Figure 6: Simulated Statistics : Maxima and Minima

Notes: These plots show the maximum and minimum of the annual occupation switching rate, and three measures of wage in equality for each simulated time series in the ensemble of 1000 simulated time series.

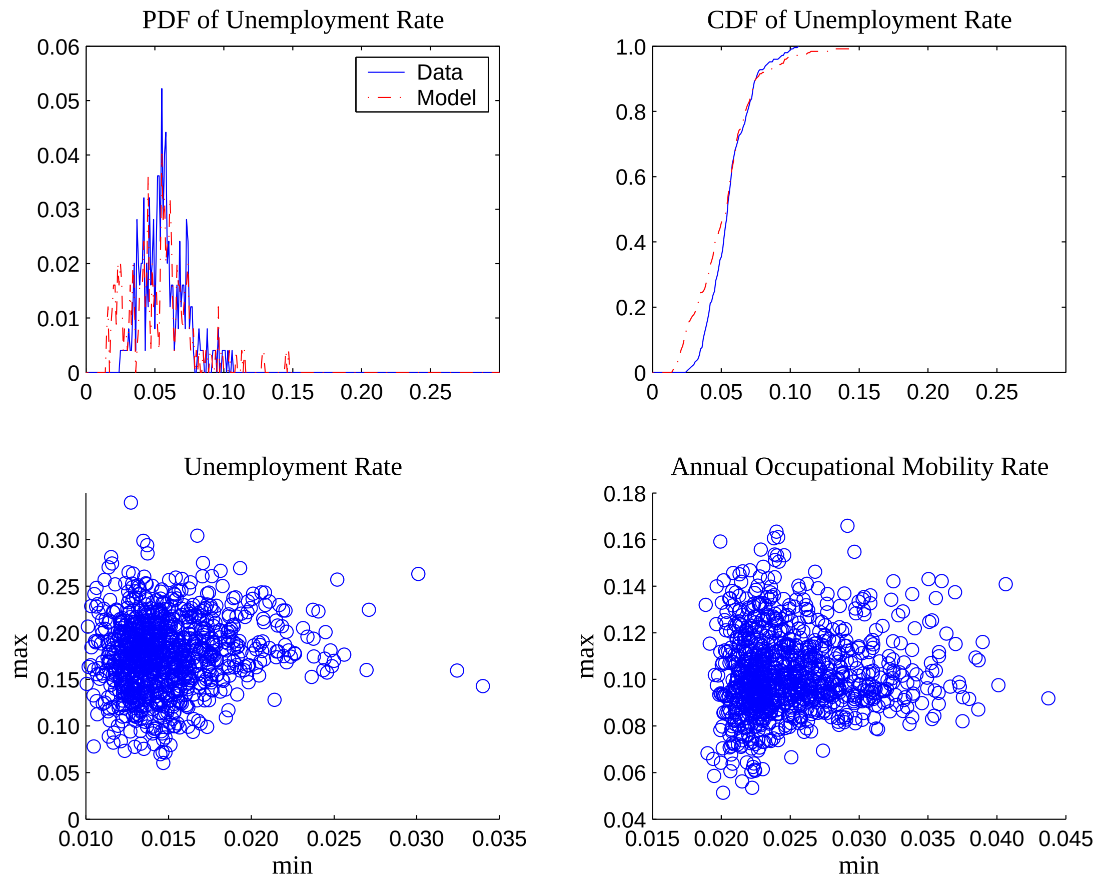

Figure 7: Time Series of Best Fit

Note : This plots pdf and cdf of the unemployment rate for the time series from the ensemble that best fits the U.S. data using a minimum mean squared deviation criterion. The histogram is constructed using bins of size 0.01 and runs from 0 to 0.3.

Table 7: Fraction of Simulations Featuring n Decades with Positive corr(LP,UR)

Number of Decades	0	1	2	3	4	5	6
$\chi = 2.75$	47.0%	31.0%	16.1%	4.4%	0.9%	0.5%	0.1%
$\chi = 3$	60.7%	25.9%	9.9%	2.8%	0.7%	0%	0%

Note : The logarithm of the data are taken and detrended using the band pass filter. For medium run frequencies we isolate frequencies between 6 and 80 quarters.

Table 8: Fraction of Simulations Featuring n Decades with Negative corr(OM,UR)

Number of Decades	0	1	2	3	4	5	6
$\chi = 2.75$	61.7%	29.3%	7.7%	1.2%	0.1%	0%	0%
$\chi = 3$	50.7%	34.4%	11.6%	2.7%	0.6%	0%	0%

Note : The logarithm of the data are taken and detrended using the band pass filter. For medium run frequencies we isolate frequencies between 6 and 80 quarters.

The ability of the model to generate qualitative sign switches in the contemporaneous correlation between the unemployment rate and average labour productivity is presented in Table 7. At medium run frequencies, 25.9% of the simulated series in the ensemble contained precisely one decade with positive correlations while 9.9% of the simulated series contained exactly two decades with positive correlations. This shows that for this particular parameterization, there is an asymmetry in the likelihood of observing decades with positive correlations between unemployment and labour productivity at medium-run. This also shows that the model can generate sign flips qualitatively as in the U.S. data.

Next we note that typically the model generates a strong positive correlation between the unemployment rate and the rate of occupational mobility. Unemployment tends to rise when a reallocation shock pushes the productivity frontier away from occupations housing a reasonable measure of workers. In a frictionless world, these workers would jump to higher productivity occupations. However, due to the idiosyncratic fixed costs of switching occupations, many of these workers must endure several periods of unemployment before they draw a cost sufficiently low to merit switching occupations. It is exactly during these periods that unemployment is high and occupational mobility increases generating the positive correlation between the variables. We can see that under the baseline parameterization, only 11.6% of the simulations in the ensemble feature two decades of negative correlation

Table 9: Fraction of Simulations Featuring n Decades with Positive corr(LP,UR) and Negative corr(OM,UR)

Number of Decades	0	1	2	3	4	5	6
$\chi = 2.75$	72.6%	22.5%	4.3%	0.6%	0%	0%	0%
$\chi = 3$	80.4%	16.5%	2.9%	0.2%	0%	0%	0%

Note : The logarithm of the data are taken and detrended using the band pass filter. For medium run frequencies we isolate frequencies between 6 and 80 quarters.

between the unemployment rate and the rate of occupational mobility. Again, there is an asymmetry in the model to generate multiple periods of negative versus positive correlations.

Next we examine the ability of the model to generate changes in wage inequality along with rises in the occupational mobility rate in a way that the two series are positively correlated. We do not have high frequency data on wage inequality so instead of comparing the model's performance to any particular set of statistics generated by the U.S. economy, we will ask whether the model can match stylized observations. Specifically, we will drop the first 8 periods and the last period from each sample run of 249 periods and then divide the remaining data into six decades as we did above. Then we calculated measures of wage inequality and occupational mobility rates at a quarterly frequency. For each decade we take the minimum and the maximum values of wage inequality and occupational mobility and calculate the percentage change. If this is large then we count that as a success for that decade. In addition, we then calculate whether the contemporaneous correlation between wage inequality and occupational mobility is sufficiently large. If there are large percentage changes in occupational mobility, wage inequality and the two variables are highly, positively correlated then we say that the model has matched the stylized observation. In order to be harsh on the model, we ask that, within a decade (as opposed to two decades) it be able to generate the within-group inequality changes as shown in Table 2 as well as the changes in occupational mobility. For a "large", positive correlation between wage inequality and occupational mobility we used values of 0.5, 0.75 and 0.9.

We can see that, given the highly volatile nature of the economy, it has little trouble generating periods in which there are large changes in both wage inequality and occupational mobility along with the two series being highly correlated.

Table 10: Fraction of Simulations Featuring n Decades of Large Changes in Mobility and Inequality with a Large, Positive Correlation

Number of Decades	0	1	2	3	4	5	6
corr(Gini,OM) $\geq$ 0.50	0.0%	0.0%	12.0%	8.9%	21.8%	40.0%	28.1%
corr(Gini,OM) $\geq$ 0.75	0.3%	3.2%	11.3%	25.7%	33.7%	19.4%	6.4%
corr(Gini,OM) $\geq$ 0.90	8.5%	23.7%	31.0%	23.6%	10.0%	3.0%	0.2%
corr($\sigma^2(\log(w))$,OM) $\geq$ 0.50	0.0%	0.9%	5.0%	15.2%	27.5%	31.9%	19.5%
corr($\sigma^2(\log(w))$,OM) $\geq$ 0.75	2.4%	8.4%	23.0%	27.9%	22.4%	12.9%	0.3%
corr($\sigma^2(\log(w))$,OM) $\geq$ 0.90	20.2%	31.5%	26.5%	14.7%	5.7%	1.3%	0.1%

Note : The logarithm of the data are taken and detrended using the band pass filter. For medium run frequencies we isolate frequencies between 6 and 80 quarters.

Table 11: Number of 20 Year Periods Featuring All Three Stylized Observations

	0	1	2	3
corr(Gini,OM) $\geq$ 0.50	82.2%	15.6%	2.2%	0.0%
corr(Gini,OM) $\geq$ 0.75	88.7%	10.5%	0.8%	0.0%
corr(Gini,OM) $\geq$ 0.90	97.5%	2.5%	0.0%	0.0%
corr($\sigma^2(\log(w))$,OM) $\geq$ 0.50	94.4%	5.5%	0.1%	0.0%
corr($\sigma^2(\log(w))$,OM) $\geq$ 0.75	98.6%	1.4%	0.0%	0.0%
corr($\sigma^2(\log(w))$,OM) $\geq$ 0.90	99.8%	0.2%	0.0%	0.0%

Lastly, in checking whether the model can produce simulations that resemble those described as stylized observations, for each sample in the ensemble we construt three twenty year periods. As above, we simply drop the first eight quarters of simulated data and the last quarter of simulated data and then divide the remaining 240 quarters into three twenty year periods. We check to see whether each of the twenty year periods features at least one decade with a positive correlation between average labour productivity and the unemployment rate, at least one decade with a negative correlation between the unemployment rate and the occupational mobility rate, and across the twenty years features large changes in mobility and inequality together with a large, positive correlation between the two. Table 11 shows these results using either the Gini coefficient or the variance of log wages as the measure of inequality. We can see that if we define a large correlation between inequality and occupational mobility as 0.5 then 15.6% of the simulations exhibit at least one twenty year period in which the three stylized observations are obtained together.

3.2.1 Distributions and Dynamics

In this subsection we provide an illustration of the response of the model economy to different initial distributions of workers across occupations. The aim is to shed some light on the main mechanism at work in the model economy. Intuitively the idea is straightforward. Workers are assigned to a particular occupation. At the beginning of each period, they obtain a fixed cost of retraining. Given the distribution of workers across occupations as well as the distribution of labour productivity across occupations, workers decide which occupation is the optimal occupation to choose if the fixed cost of retraining is incurred. Knowing that all other workers face the same choice, each worker then decides whether it is best to retrain or not. At the same time, if retraining is not optimal, then the worker will either be rest unemployed or employed at the current wage offered in his/her current occupation. The complexity of the problem arises from the fact that the distribution of wages expected to be paid in each occupation depends on the distribution of workers across the occupations. As workers care about the stream of wages they will earn over their lifetime, they will prefer to move to high productivity occupations but also occupations not over-supplied with workers. They also consider the continuation value of being in any particular occupation which means they consider the likely future productivity levels of each occupation and the expected supply of workers at each occupation.

When intermediate goods are perfect substitutes in the production of the consumption good, there is no feedback between the distribution of workers and the wage paid. In such a partial equilibrium world, wages are pinned down at the level of output per worker in each occupation and the distribution of workers does not act to smooth out wages across occupations. This results in bunching of workers in occupations so when many workers are bunched in highly productive occupations and there is a reallocative shock, the unemployment rate spikes. We provide an example of the partial equilibrium behaviour of this model in a following section of the paper.

To illustrate how the distribution of workers matters, Figures 8 and 9 take two distributions from a single sample run of 249 periods. We conduct two exercises, labeled Simulation #1 and Simulation #2. In each of these exercises we take the initial distribution of workers as one that is realized from the sample run (so that we know these distributions are feasible in equilibrium) and run the counterfactual by shutting down the reallocation shocks. We can see from Figure 8 that in Simulation #1, the initial distribution has a reasonable mass of workers in the occupations that are in the intervals $\delta \in [-3,-2]$. Over time, workers in this interval switch to occupations in the vicinity of $\delta = -1.5$ and in doing so reduce the amount of wage inequality in the economy. A similar story occurs in the inverval $\delta \in [2,3]$ as these workers switch to occupations around $\delta = 1.5$. The dynamics over time generated by this occupation switching is described by the occupational mobility rate falling over time along with the fall in wage inequality. In this simulation, the variance of the initial distribution of workers across occupations is high and so there is a large measure of workers who would prefer to move to higher paying observations. Over time, as these workers draw low fixed costs of retraining, they move to higher paying jobs which lowers the wage inequality. Furthermore, as more and more workers settle into occupations for which it is not worth incurring fixed costs to retrain, the rate of occupational mobility falls. This provides an explanation for the relation of falling occupational mobility and wage inequality over the simulation in tandem with rising average labour productivity and lower unemployment rates.

The dynamics in Simulation #2 are shown in Figure 9. We can see that in this example that the initial distribution is such that most of the workers are in occupations near the productivity frontier. A small measure of workers is located in the occupations $\delta \in [-\pi,-3] \cup [3,\pi]$. The initial unemployment rate is low; close to 1.5% and the rise over time is almost totally accounted for by new labour force entrants being exogenously dropped into these low productivity occupations.

3.3 A Partial Equilibrium Parameterization

In this section we show the behaviour of the model in a partial equilibrium setting in which the distribution of workers across occupations does not affect the distribution of wages across occupations. On the one hand, the example shows that the property that the contemporaneous correlation between the unemployment rate and average labour productivity can switch signs over a simulation of 249 periods is not unique to the general equilibrium setting. What matters for the results is that the distribution of workers across occupations at the time that a reallocation shock hits the economy. However, in the partial equilibrium

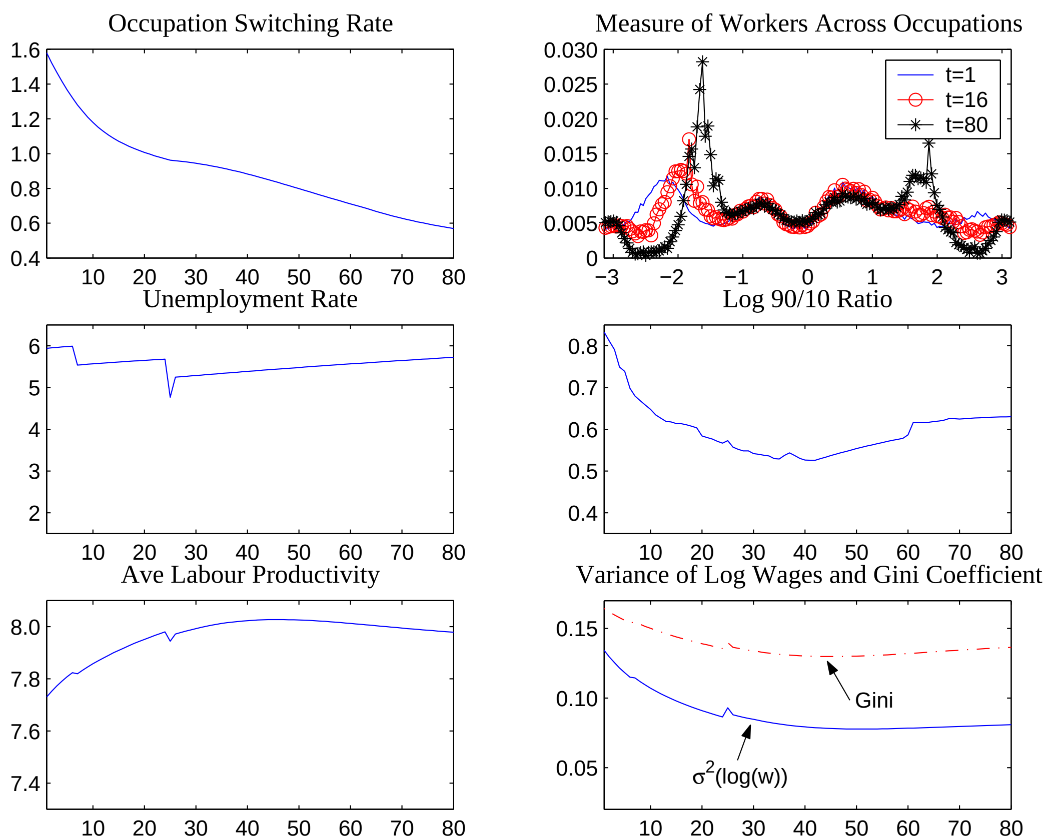

Figure 8: Simulated Impulse Responses

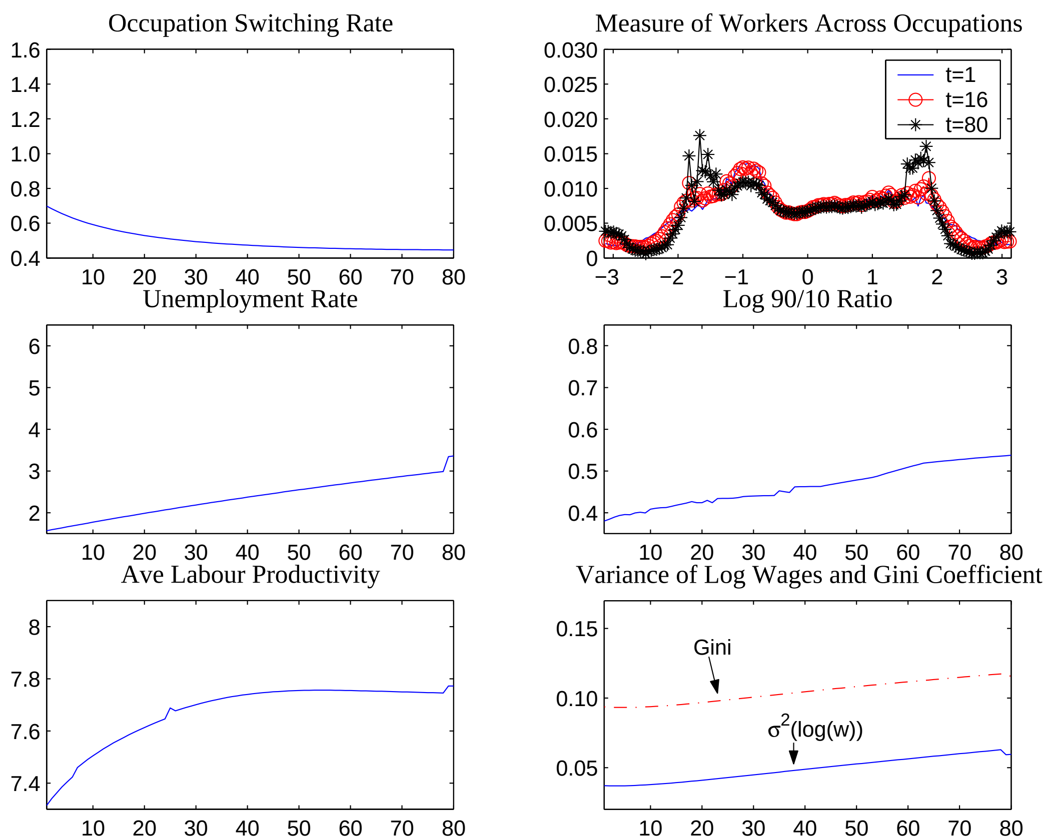

Figure 9: Simulated Impulse Responses

Table 12: Parameter Values

Parameter	P.E.	
ξ	1.75	Shape parameter for labour productivity function $g(\delta)$
α_ϵ	120	Parameter governing distribution of shocks
μ_z	1.65	Mean of idiosyncratic shocks from Gamma Distribution
σ_z^2	0.90	Variance of idiosyncratic shocks from Gamma Distribution
b	0.29	Fraction of $\overline{a}$ enjoyed by unemployed
η	40	Weight on quadratic retraining cost

setting the frequency with which negative correlations between the unemployment rate and occupational mobility rate arise are reduced drastically. Thus it appears as though having general equilibrium effects to spread out the distribution of workers across occupations plays a large role in this property of the model.

Under the partial equilibrium parameterization, we let intermediate goods be perfect substitutes in the final good aggregator. We adjust the parameters so that the economy is in the vicinity of the targets outlined for the general equilibrium case. For the partial equilibrium example, the values of the parameters used are listed in Table 12. Given the parameters, for a simulation of 50000 periods there is an average unemployment rate of 5.3%, an annual occupational switching rate of 5.5%, a replacement ratio 34% and a median distance moved of 0.103. These are all within reasonable distance of four of our targeted statistics for the U.S. economy. The remaining targets are the volatility of labour productivity and the first-order autocorrelation of labour productivity. We are able to obtain a parameterization where the observed volatility and autocorrelation of labour productivity from the U.S. economy lies within the 95% bands of the simulated counterparts from an ensemble of simulated time series.

Comparing the correlations in Table 13 to the baseline results in Table 5 we can see that the volatility, persistence and correlation between the unemployment rate and average labour productivity do not vary much across the two parameter specifications. By design the volatility and persistence of average labour productivity across the two parameterizations are similar as these are two of the targeted statistics in choosing the parameter values. Note that between the partial equilibrium setting and the general equilibrium setting that the volatility of occupational mobility is lower in the partial equilibrium setting with the

Table 13: Partial Equilibrium Model : Simulated Statistics

	u	Y/e	n
Standard Deviation	0.3155	0.0226	0.1757
	(0.2207,0.4073)	(0.0138,0.0337)	(0.1073,0.2553)
Quarterly Autocorrelation	0.9438	0.9486	0.9410
	(0.9055,0.9685)	(0.9105,0.9736)	(0.9029,0.9680)
Correlation Matrix			
u	1.0000	-0.6732	0.8144
		(-0.8962,-0.3124)	(0.6522,0.9161)
$\frac{Y}{e}$	-	1.0000	-0.8815
			(-0.9709,-0.7089)
n	-	-	1.0000

Notes: To construct the data the model is simulated 1000 times. The logarithm of the simulated data are taken and detrended using the band pass filter isolating frequencies between 6 and 80 quarters. The mean of the relevant statistics are presented in the table with their standard deviations in parenthesis.

correlation between the unemployment rate and occupational mobility becoming much more positive; to the point that the ensemble mean in the general equilibrium setting lying outside the 95% bounds of the partial equilibrium setting. This plays a important role in establishing the difference between the performance of the partial equilibrium model and the general equilibrium model as discussed later in this section. Also the correlation between average labour productivity and occupational mobility becomes more negative relative to the baseline general equilibrium setting.

Stark differences between the partial equilibrium and general equilibrium settings arise when examining the statistics on wage inequality. Across all three measures of wage inequality, there is a dramatic reduction of wage inequality in the partial equilibrium setting. Some of this comes from the need to spread productivity out across occupations in the P.E. setting in order to hit the targeted statistics from the U.S. data. Also, as workers are paid their labour productivity in the partial equilibrium model, if workers bunch up in a small set of occupations near the productivity frontier then wage inequality will typically be low. As the distribution of workers across occupations does not affect wages, the incentives of the individual worker to enter an occupation near the frontier is not reduced if many workers are already attached to that occupation. Given the targets for the median distance of occupation switches and the frequency of switches, we observe much more bunching under the partial equilibrium model. Across runs of 50000 periods, the mean of the variance of workers across occupations is 2.52 under the partial equilibrium parameterization and 2.59 under the general equilibrium setting. However, the standard deviation of the variance of workers across occupations is 0.57 in the partial equilibrium model and 0.43 under the general equilibrium setting showing that there are more ups and downs in the variance of the distribution of workers across occupations in the partial equilibrium model.

It can also be seen that all three measures of wage inequality become highly positively correlated with the rate of occupational mobility with the ensemble mean of the G.E. setting falling below the 95% bounds of the P.E. setting for all three correlations. With respect to the correlations between unemployment rates and wage inequality, we see that in the P.E. setting that the log 90/10 ratio and the Gini coefficient all become much more positive relative to the G.E. setting and the variance of log wages becomes much less positively correlated. Also, the correlation between wage inequality and occupational mobility increases. As the wages paid across occupations relative to the frontier occupation is invariant to the distribution of workers across occupations, when there is a reallocation shock unemployment rises as does the amount of occupation switching. When many workers switch occupations wage inequality measures rise because some workers obtain higher wages in comparison to those who were in the same occupation pre-shock but have yet to draw low idiosyncratic fixed costs of switching. In the absence of general equilibrium effects, wages are not depressed in occupations experiencing an inflow of workers and wages are not buoyed in occupations experiencing an outflow of workers. This acts to magnify the positive correlation between occupational mobility rates and measures of wage inequality.

The first row of Table 15 displays the fraction of samples in the ensemble of 1000 simulations for which the partial equilibrium model displays n decades with positive contemporaneous correlations between the unemployment rate and average labour productivity. After filtering the simulated data to isolate medium run frequencies the partial equilibrium model produces sign switches in the contemporaneous correlation. Furthermore, it is more likely to observe sample runs with less than two decades exhibiting contemporaneous correlations with only one decade of positive correlation occurring between unemployment and labour

Table 14: Partial Equilibrium Model : Wage Dispersion Measures

	$\sigma^2(\log(w))$	$\log\left(\frac{w_{90}}{w_{10}}\right)$	Gini
Mean	0.0725	0.6408	0.1259
	(0.0558,0.0864)	(0.5066,0.7468)	(0.1046,0.1432)
Standard Deviation	0.1813	0.1396	0.1113
	(0.1246,0.2425)	(0.0976,0.1866)	(0.0736,0.1548)
Quarterly Autocorrelation	0.9363	0.9378	0.9418
	(0.8977,0.9632)	(0.8999,0.9635)	(0.9045,0.9679)
Correlation Matrix	$\sigma^2(\log(w))$	$\log\left(\frac{w_{90}}{w_{10}}\right)$	Gini
u	0.8813	0.8918	0.8559
	(0.7813,0.9417)	(0.8098,0.9468)	(0.7287,0.9328)
$\frac{Y}{e}$	-0.6977	-0.7611	-0.8439
	(-0.9247,-0.2877)	(-0.9476,-0.4173)	(-0.9714,-0.5835)
n	0.9271	0.9533	0.9879
	(0.8195,0.9805)	(0.8739,0.9894)	(0.9621,0.9985)

Notes: To construct the data the model is simulated 1000 times. The logarithm of the simulated data are taken and detrended using the band pass filter isolating frequencies between 6 and 80 quarters. The mean of the relevant statistics are presented in the table with their standard deviations in parenthesis. The exception to this are the statistics for the mean for which we report the mean of the unfiltered data (without taking the logarithm) generated by the model economy.

Table 15: Partial Equilibrium Model

Number of Decades	0	1	2	3	4	5	6
$\sigma_{UR,LP} > 0$	55.2%	34.7%	8.4%	1.4%	0.3%	0.0%	0.0%
$\sigma_{UR,OM} < 0$	95.3%	4.7%	0.0%	0.0%	0.0%	0.0%	0.0%
$\sigma_{UR,LP} > 0$, $\sigma_{UR,OM} < 0$	96.9%	3.1%	0.0%	0.0%	0.0%	0.0%	0.0%

productivity being most likely if any sign switches occur at all. The intuition behind this finding is the same as under the general equilibrium model.

The partial equilibrium model has a much harder time producing decades in which the unemployment rate and occupational mobility co-move negatively at medium run frequencies. Particularly, only 4.7% of the simulations in the ensemble produce one decade in which the unemployment rate and the occupational mobility rate produce a negative correlation. Clearly from Table 13 it can be seen that across the simulations comprising the ensemble, the average mean correlation between the unemployment rate and the occupational mobility rate stands at 0.81. In contrast the G.E. model produces an average correlation across the ensemble of 0.59 and the lower end of the 95% band is 0.19 which is considerably lower than 0.65 as produced by the P.E. model. Notably, the average correlation from the G.E. model sit below the 95% bound of the P.E. model.

As noted in the discussion above, without the general equilibrium effects allowing the distribution of workers to feedback into wages, the individual incentives to switch occupations is not affected by the measure of workers attached to any particular occupation. In order to obtain an increase in the unemployment rate and a simultaneous fall in the rate of occupational mobility it needs to be the case that a reallocation shock adversely affects occupations in which many workers are attached and the measure of workers who choose not to switch occupations rises relative to previous periods. This may arise in a general equilibrium model if the occupations "close-by" are already housing many workers so that wages in those occupations would be expected to stay low for many periods. This effect is missing the in the partial equilibrium model. Thus for the rise in the unemployment rate to be accompanied by a fall in the rate of occupational mobility, it must be that the sequence of shocks is such that one shock pushes workers to move into low paying occupations and

Table 16: Fraction of Simulations Featuring n Decades of Large Changes in Mobility and Inequality with a Large, Positive Correlation

Number of Decades	0	1	2	3	4	5	6
corr(Gini,OM) $\geq$ 0.50	0.0%	0.0%	0.0%	0.3%	0.3%	17.8%	78.5%
corr(Gini,OM) $\geq$ 0.75	0.0%	0.0%	0.0%	0.4%	4.2%	20.7%	74.7%
corr(Gini,OM) $\geq$ 0.90	0.0%	0.0%	0.0%	0.1%	6.2%	27.2%	65.5%
corr($\sigma^2(\log(w))$,OM) $\geq$ 0.50	0.0%	0.0%	0.1%	1.0%	8.1%	29.8%	61.0%
corr($\sigma^2(\log(w))$,OM) $\geq$ 0.75	0.0%	0.1%	0.7%	3.3%	16.5%	36.8%	42.6%
corr($\sigma^2(\log(w))$,OM) $\geq$ 0.90	0.1%	1.4%	5.2%	16.1%	30.3%	32.0%	14.9%

Note : The logarithm of the data are taken and detrended using the band pass filter. For medium run frequencies we isolate frequencies between 6 and 80 quarters.

Table 17: P.E. : Number of 20 Year Periods Featuring All Three Stylized Observations

	0	1	2	3
corr(Gini,OM) $\geq$ 0.50	97.0%	3.0%	0.0%	0.0%
corr(Gini,OM) $\geq$ 0.75	97.1%	2.9%	0.0%	0.0%
corr(Gini,OM) $\geq$ 0.90	97.7%	2.3%	0.0%	0.0%
corr($\sigma^2(\log(w))$,OM) $\geq$ 0.50	97.3%	2.7%	0.0%	0.0%
corr($\sigma^2(\log(w))$,OM) $\geq$ 0.75	98.0%	2.0%	0.0%	0.0%
corr($\sigma^2(\log(w))$,OM) $\geq$ 0.90	99.4%	0.6%	0.0%	0.0%

then another shock leaves this large mass of workers in a "rest-unemployed" occupation. However given the calibrated frequency of occupation switching and calibrated median distance, when workers are near the rest-unemployed occupations they will try to switch to higher paying occupations and when they are able to switch they move quite a distance away from these undesirable occupations. This means that it will not often be the case that many workers get trapped in rest-unemployed or low wage occupations in which few are able to move out of when a bad shock hits.

As with the general equilibrium model, the partial equilibrium model features a tremendous amount of volatility in the occupation mobility rate. Thus it is fairly easy for the model to produce periods during which there are large percentage changes in the rate of occupational mobility. From Table 17 we can see that the impediment for the partial equilibrium model to matching all three stylized observations simultaneously is the inability to produce periods in which the unemployment rate and the occupational mobility rate are negatively correlated.

4 Discussion

One issue in dealing with the interpretation of the model is in deciding whether workers in the mobility state should be considered as unemployed. In our quantitative work we have assumed that these workers are non-participants and so the level of unemployment is purely a measure of rest unemployed workers. One reason for leaving workers in the mobility state out of the labour force is that during the period of occupation switching they are not available for work. These workers make the choice to leave the labour force for a period to retrain. One example of this would be that workers who switch occupations become full-time students in a vocational training program. This may be an objectionable assumption when calibrating the model and an option would be to consider those who are switching occupations as search unemployed.

Another issue that merits investigation is the dynamic relationship of labour productivity across occupations. We have assumed that the relative productivity of occupations that are a particular distance from the most productive occupation is time invariant, i.e. the $g(\delta)$ function is fixed. While this is a simplification, the shape of the distribution need not be made time invariant so long as the shape can be indexed by a small number of state variables. This would increase computational time but allow more flexibility in addressing the data. In addition, to add flexibility in addressing the data, occupations could be given mean zero shocks every period as long as these shocks are i.i.d. and independent over time. For other applications the $g(\delta)$ function can be used to model shifts in demand for the output across occupations. Nonetheless, our modeling assumption requires some level of systematic comovement of productivity between the most productive occupation and other occupations which should be examined carefully if the data becomes available.

A margin of the model for which we have yet to carry out extensive investigation is the implication on unemployment duration and wages. Specifically, we have not investigated whether workers who are displaced, and spend an uninterrupted spell as rest unemployed experience lower wages upon re-employment relative to their wage the period before displacement. This could be contrasted with the experience of those who are displaced but choose to retrain and find employment in new occupations. Our model lacks an important dimension to address this question in that workers do not accumulate occupational (or task) capital which can be destroyed upon switching occupations. However, as our model does consider distance between occupations, it may serve as a building block for incorporating the idea that, while workers may switch occupations, they do direct their choice of new occupations so as to minimize their loss of task-capital.

5 Conclusion

This paper presents a dynamic model of occupational choice in which an economy is subjected to aggregate reallocation shocks and workers may choose to incur costs to retrain in order to move to occupations that pay higher wages. Displaced workers who temporarily face high fixed costs of switching occupations suffer through rest unemployment spells. It is shown that the delicate interaction between changes in the distribution of workers across occupations and the distribution of labour productivity across occupations generates qualitative switches in the contemporaneous correlation between the unemployment rate and average labour productivity; a property that has emerged in recent decades in the U.S. economy. Importantly, the model obtains the results without relying on structural shifts in any parameters and only requires a single type of shock. We also showed that the model can generate periods in which the unemployment rate and the rate of occupational mobility is negatively correlated. Lastly, we showed that the model economy can be quite volatile and generates periods in which occupational mobility and wage dispersion are positively correlated while experiencing large changes.

We hope that this model can be used as a building block in order to examine issues of optimal unemployment insurance versus retraining subsidies. However, this will likely require additional computing power and is left as future work.

Appendix

A Solution Method

The solution method used in this paper is similar to the methods proposed in Reiter (2002) and Reiter (2010).

A.1 Outline of the Algorithm

In order to approximate the equilibrium distribution of workers across occupations, we use a grid of moments $\mathcal{M}$ which is formed as a Cartesian product from the set of one-dimensional grids for a set of m moments, $\mathcal{M}_j$ with $j \in \{1, ...m\}$. In our particular application we used two moments, the mean and the variance of the location of workers across a set of bins in the interval $[-\pi, \pi]$. Let $\hat{m}$ denote an element of $\mathcal{M}$. We partitioned the interval $[-\pi, \pi]$ into n_δ bins so that the density of workers can be approximated by a histogram following the procedure of Young (2010). Denote the set of mid-points of the bins by $\Delta = \{\delta_1, ..., \delta_{n_\delta}\}$.

1. For each grid point in $\hat{m} \in \mathcal{M}$ guess at the histogram $\psi(\delta, \hat{m})$, for all $\delta \in \Delta$.

   (a) For each grid point $\hat{m} \in \mathcal{M}$ guess at the transition function $\hat{\Gamma} : [-\pi, \pi] \times \mathbb{R}_+ \to [-\pi, \pi] \times \mathbb{R}_+$ where, as we use the mean and variance, we have $\hat{m} \in [-\pi, \pi] \times \mathbb{R}_+$. This provides us with a guess $\hat{m}' = \Gamma(\hat{m})$ for each $\hat{m} \in \mathcal{M}$.

      i. Taking the functions $\psi(\delta, \hat{m})$ and $\hat{m}' = \Gamma(\hat{m})$ as given solve for the individual decision rules using the value function iteration procedure discussed below. Denote the optimal decision rule for choosing the employment state by $\iota(\delta, \hat{m})$ and the decision rule for retraining by $x(\delta, \hat{m})$.

   (b) For each grid point $\hat{m} \in \mathcal{M}$, using the decision rules $\iota(\delta, \hat{m})$ and $x(\delta, \hat{m})$ along with the histogram $\psi(\delta, \hat{m})$ calculate the $\tilde{m}'$ consistent with individual optimization and the conjectured equilibrium distribution. Check to see if $\tilde{m}' = \hat{m}'$. If so, we have found the individual rules and the equilibrium transition function that are consistent with each other taking the equilibrium distribution as given.

2. Taking the transition function and decision rules from above, simulate the aggregate economy for T periods saving the distribution of workers for each $t = 1, ..., T$. Now using the procedure below, approximate the equilibrium density function $\tilde{\psi}(\delta, \hat{m})$. For each $\hat{m} \in \mathcal{M}$ check to see if the distance between $\psi(\delta, \hat{m})$ and $\tilde{\psi}(\delta, \hat{m})$ is sufficiently small. If so, then stop. Otherwise, repeat until convergence is achieved.

A.2 Value Function Iteration

In order to solve the individual worker's dynamic programming problem, we take as given the transition function for the moments as well as the distribution of workers for each $\hat{m} \in \mathcal{M}$. The value functions for the workers have three arguments, δ, μ and σ^2. We solve the dynamic programming problem with a slight alteration of the method proposed by Barillas and Fernández-Villaverde (2007). We do not use information on the derivative of the value functions as is standard in the endogenous grid point method because of problems that potential kinks in the continuation values of employment, unemployment or retraining pose along with the discrete choice nature of the worker's problem. Particularly, by using derivatives of the continuation value, due to kinks, it is possible that two occupations, say $\hat{\delta}' = \delta + \hat{x}$ and $\tilde{\delta}' = \delta + \tilde{x}$ are optimal continuation occupations from the same current occupation, δ even though, one of the occupations may never be an optimal occupation to move to from δ in equilibrium. The assumption, by using the derivative of the continuation occupation, in solving backwards as in the endogenous grid point method (EGM) of Carroll (2006) and Barillas and Fernández-Villaverde (2007) is that every occupation in the continuation grid of occupations can be reached in equilibrium from some occupation in the grid of current occupations; an assumption that is not necessarily true in the equilibrium of our model.

Thus, instead of exploiting derivatives, we use the following procedure which combines some ideas of traditional value function iteration procedures and some ideas from the EGM methods. Particularly, we solve the dynamic programming problem forwards (in contrast to backwards as in the EGM) which allows us to avoid the problem detailed in the previous paragraph. We follow the EGM method in that we reduce the number of times that we solve for the expected continuation values per iteration of the value functions which reduces computational time immensely.[17]

We now describe the procedure. Define a grid on the continuation occupation $G' \equiv \{\delta'_1, ..., \delta'_{N'}\}$ and a grid for current occupations $G \equiv \{\delta_1, ..., \delta_N\}$. Construct a large (symmetric) grid for values of x, $X \equiv \{x_1, ..., x_{N_x}\}$ which contains values between $-\pi$ and π. The algorithm to iterate on the value functions is

1. Start with $k = 0$ and guess values for $V^k(\delta, \mu, \sigma^2)$, $U^k(\delta, \mu, \sigma^2)$, and $T^k(\delta, \mu, \sigma^2)$.

2. Using Gaussian quadratures, for each $(\hat{\delta}, \mu, \sigma^2)$, where $\hat{\delta}$ denotes the end of period occupation for a worker, construct the expected continuation values

$$W^k(\hat{\delta}, \mu, \sigma^2) = \int\int \max\left\{T^k(\delta', \mu', \sigma^{2'}), U^k(\delta', \mu', \sigma^{2'}), V^k(\delta', \mu', \sigma^{2'})\right\} dF(\delta'|\hat{\delta})dH(z')$$

3. Construct a grid $G(\delta) = \{\delta + x_1, ..., \delta + x_{N_x}\}$ for each δ and then interpolate the vector $W(\hat{\delta}, \mu, \sigma^2)$ across this grid. This yields the expected continuation value for each end of the period occupation in the grid $G(\delta)$. Then for each δ, it is simple to construct the current period value functions $V^k(\delta, \mu, \sigma^2)$, $U^k(\delta, \mu, \sigma^2)$, and $T^k(\delta, \mu, \sigma^2)$. For each $\delta \in G$ construct the functions

$$\begin{aligned} V^{k+1}(\delta, \mu, \sigma^2) &= u(w(\delta, \mu, \sigma^2) + \beta W^k(\delta, \mu, \sigma^2) \\ U^{k+1}(\delta, \mu, \sigma^2) &= u(b) + \beta W^k(\delta, \mu, \sigma^2) \\ T^{k+1}(\delta, \mu, \sigma^2) &= \max_{x \in X}\left\{u(b) - \varphi(x) + \beta W^k(\delta + x, \mu, \sigma^2)\right\}. \end{aligned}$$

4. If

$$\sup_{i,j,N}\left\{\frac{\left\|\begin{bmatrix} V^{k+1}(\delta_N, \mu_i, \sigma_j^2) - V^k(\delta_N, \mu_i, \sigma_j^2) \\ U^{k+1}(\delta_N, \mu_i, \sigma_j^2) - U^k(\delta_N, \mu_i, \sigma_j^2) \\ T^{k+1}(\delta_N, \mu_i, \sigma_j^2) - T^k(\delta_N, \mu_i, \sigma_j^2) \end{bmatrix}\right\|}{1 + \sup_{i,j,N}\left\|\begin{bmatrix} V^k(\delta_N, \mu_i, \sigma_j^2) \\ U^k(\delta_N, \mu_i, \sigma_j^2) \\ T^k(\delta_N, \mu_i, \sigma_j^2) \end{bmatrix}\right\|}\right\} \geq 1.0e^{-6}$$

[17] As a check of this procedure, we solved Model 1 of Barillas and Fernández-Villaverde (2007), found that the solutions were virtually identical and that our procedures was slightly slower. However, our procedure was able to solve certain cases in our model which ran into difficulties when using deriviatives of the continuation value functions.

then $k \rightsquigarrow k+1$ and go to 2. Otherwise, stop and store the value functions, $V(\delta, \mu, \sigma^2)$, $U(\delta, \mu, \sigma^2)$, $T(\delta, \mu, \sigma^2)$ and the decision rule $X(\delta, \mu, \sigma^2)$.

A.3 Simulating the Economy

In order to simulate the economy we use a method similar to that proposed by Young (2010). Specifically, we use a histogram over a fixed grid of the occupations to construct moments for the next period. Suppose we break the interval $[-\pi, \pi]$ into N_B equal length bins whose midpoints serve as the "location" of the bin. Once we know the measure of workers in each bin, we can easily calculate the moments desired. In simulating the model for T periods we then use the following procedure :

1. Take the histogram of workers at the beginning of the period t as given.

2. Given this histogram, construct the mean and variance of the distribution. Denote the period t mean by μ_t and the variance by σ_t^2.

3. Use the value functions as obtained from the value function iteration and interpolate across the moments in $\mathcal{M}$ to obtain $V(\delta, \mu_t, \sigma_t^2)$, $U(\delta, \mu_t, \sigma_t^2)$, $T(\delta, \mu_t, \sigma_t^2)$ and $X(\delta, \mu_t, \sigma_t^2)$ for all $\delta \in G$.

4. Using these value functions, for each bin, determine the cut-off cost (from the distribution of idiosyncratic fixed cost of retraining) for which workers would prefer to retrain rather than be employed or unemployed. Also determine, for each bin, if workers are not to retrain whether they would prefer to be employed or unemployed.

5. For those who elect to retrain, determine the bin in which their new occupation falls. Assign all workers who retrain to their new bins.

6. Draw an aggregate shock for period $t+1$. Adjust the bins for all workers given the end of period t histogram and the period $t+1$ reallocation shock.

7. For each bin, let a fraction ρ of the workers withdraw from the labour force. In each bin add a measure $\frac{(1-\rho)}{N_B}$ of new workers.

8. If $t <= T$ then stop. Otherwise, $t \rightsquigarrow t+1$ and go to 1.

A.4 The Transition Function for the Moments

In order to determine the transition function for the moments we follow the steps below. Suppose N_ϵ nodes are used to evaluate the Gaussian quadrature in the distribution for the reallocation shocks, that there are N_μ elements in the grid for the mean and that there are N_{σ^2} elements in the grid for the variance.

1. For possible triple $(\mu_i, \sigma_j^2, \epsilon_k)$, conjecture the continuation moment pair $\{\mu'_{ijk}, \sigma^{2'}_{ijk}\}$.

2. Using these continuation moments, for each $\{\mu, \sigma^2\} \in \mathcal{M}$, construct the continuation values

$$W(\delta, \mu, \sigma^2) = \int \int \max \left\{ T(\delta', \mu', \sigma^{2'}), U(\delta', \mu', \sigma^{2'}), V(\delta', \mu', \sigma^{2'}) \right\} dF(\delta'|\delta) dH(z')$$

   and then construct the functions

$$\begin{aligned} \hat{V}(\delta, \mu, \sigma^2) &= u(w(\delta, \mu, \sigma^2) + \beta W(\delta, \mu, \sigma^2) \\ \hat{U}(\delta, \mu, \sigma^2) &= u(b) + \beta W(\delta, \mu, \sigma^2) \\ \hat{T}(\delta, \mu, \sigma^2) &= \max_{x \in X} \left\{ u(b) - \varphi(x) + \beta W(\delta + x, \mu, \sigma^2) \right\}. \end{aligned}$$

3. For each $\{\mu, \sigma^2\} \in \mathcal{M}$, simulate the economy for one period and find the realized moment pairs $\{\tilde{\mu}'_{ijk}, \tilde{\sigma}^{2'}_{ijk}\}$. Note that this is a *realized* pair meaning one pair for each possible shock $\epsilon' \in \{\epsilon_1, ..., \epsilon_{N_\epsilon}\}$.

4. For some set of weights $\{\omega_{i,j,k}\}$, if

$$\sup \left\{ \frac{\left| \omega_{i,j,k} \begin{bmatrix} \tilde{\mu}_{ijk} - \mu'_{ijk} \\ \tilde{\sigma}^{2'}_{ijk} - \sigma^{2'}_{ijk} \end{bmatrix} \right|}{1 + \sup \left| \omega_{i,j,k} \begin{bmatrix} \tilde{\mu}_{ijk} - \mu'_{ijk} \\ \tilde{\sigma}^{2'}_{ijk} - \sigma^{2'}_{ijk} \end{bmatrix} \right|} \right\} \geq 1.0e^{-4}$$

   then set $[\mu'_{ijk}, \sigma^{2'}_{ijk}] = \lambda[\mu'_{ijk}, \sigma^{2'}_{ijk}] + (1 - \lambda)[\tilde{\mu}'_{ijk}, \tilde{\sigma}^{2'}_{ijk}]$, where $\lambda \in [0, 1)$ and go to 2. Otherwise, stop.

In this process we introduce a set of weights $\{\omega_{i,j,k}\}$ in order to reduce the computation time. Effectively, the weights may be chosen to down-weight the importance of nodes that are unlikely to be reached in equilibrium, for example, for mean and variances that are at the extremes of their respective grids or for reallocation shocks that are of very low probability of being drawn.

A.5 The Equilibrium Distribution

In order to characterize the equilibrium distribution, necessary to determine the equilibrium wages at each triplet (δ, μ, σ^2) for each $\delta \in G$, $\mu \in \mathcal{M}_1$ and $\sigma^2 \in \mathcal{M}_2$ we follow a procedure set forth by Reiter (2002) and Reiter (2010). We repeat his procedure here. First we construct what is referred to as a "reference distribution" for each $[\mu, \sigma^2] \in \mathcal{M}_1 \times \mathcal{M}_2$.

1. Initialize the reference distribution $\psi^R(\delta, \mu, \sigma^2)$ either by a uniform distribution or by some other means.[18]

2. Starting with this reference distribution, simulate the economy for $T_{Burn} + T$ periods and then discard the statistics from the first T_{Burn} "burn-in" periods. Store the mean, variance and histogram from each of these T periods. Let $\psi_{sim,t}$ denote the simulated histogram from period $t = 1, ..., T$ (it is a vector), $\mu_{sim,t}$ and $\sigma^2_{sim,t}$ be the corresponding mean and variance from period t of the stored simulation data.

3. Update the $\psi^R(\delta, \mu, \sigma^2)$ by the weighted sum

$$\begin{aligned} \psi^R(\delta, \mu, \sigma^2) &\equiv \sum_{t=1}^{T} \gamma_t \psi_{sim,t} \\ \gamma_t &\equiv \frac{d([\mu, \sigma^2], [\mu_{sim,t}, \sigma^2_{sim,t}])^\nu}{\sum_{t=1}^{T} d([\mu, \sigma^2], [\mu_{sim,t}, \sigma^2_{sim,t}])^\nu} \end{aligned}$$

   where, in our application, $d(\cdot)$ was the Euclidean metric and the weight ν was chosen to be -4.

Given the reference distributions, $\psi^R(\delta, \mu, \sigma^2)$ for each pair $(\mu, \sigma^2) \in \mathcal{M}_1 \times \mathcal{M}_2$, we then constructed the "proxy distribution". Importantly, note that the mean and variance of the reference distributions are not equal to μ and σ^2, respectively. In order to adjust

[18] In our computations, we simulate the partial equilibrium model in order to obtain some realizations for equilibrium distribution given a set of parameters. Then we solved our general equilibrium model many times while reducing the elasticity parameter χ each time. This was the most time consuming process in calibrating our model.

each of these reference distributions we used the following steps. For each (μ_i, σ_j^2) pair, for $i = 1, ..., N_\mu$, $j = 1, ..., N_{\sigma^2}$, we calculated the mean and variance of the reference distribution $\psi(\delta, \mu_i, \sigma_j^2)$. Denote this reference pair as $(\mu_{R,i}, \sigma_{R,j}^2)$.

As we used a histogram to approximate the continuous distriubtion, we use the notation ψ_l to denote the l^{th} bin in the histogram with $l = 1, ..., N_\delta$. Then we solved the following constrained optimization problem for each $(\mu, \sigma^2) \in \mathcal{M}$.

$$\min_{\{\psi_l\}_{l=1}^{N_\delta}} \sum_{l=1}^{N_\delta} (\psi_l - \psi_l^R)^2$$

subject to the constraints $\psi_l \geq \frac{1-\rho}{N_\delta}$ and

$$\begin{bmatrix} \delta_1 & \delta_2 & ... & \delta_{N_\delta} \\ \delta_1^2 & \delta_2^2 & ... & \delta_{N_\delta}^2 \\ 1 & 1 & ... & 1 \end{bmatrix} \begin{bmatrix} \psi_1 \\ \psi_2 \\ \vdots \\ \psi_{N_\delta} \end{bmatrix} = \begin{bmatrix} \mu \\ \sigma^2 + \mu^2 \\ 1 \end{bmatrix}.$$

Once this step is complete we have a set of proxy distributions $\psi(\delta, \mu, \sigma^2)$ which we use as approximations of the equilibrium distributions.

A.5.1 Updating the Proxy Distributions

In order to update the proxy distributions so that they are consistent with equilibrium behaviour, we guess a set of proxy distributions $\psi(\delta, \mu, \sigma^2)$ for each pair $(\mu, \sigma^2) \in \mathcal{M}$ to be the equilibrium distributions. Using these guesses we solve for the individual value functions and decision rules as well as the transition function for the aggregate moments. With the individual decision rules in hand, we simulate the economy for $T_0 + T$ periods and then drop the data from the first T_0 periods. Using the histograms from the simulation we then construct updates for the proxy distribution $\tilde{\psi}(\delta, \mu, \sigma^2)$. We then construct deciles from each of the original proxy distribution, denote these by $D(\mu, \sigma^2)$ and also from the updated proxy distributions, denoted by $\tilde{D}(\mu, \sigma^2)$. Then for some set of weights $\{\omega_{i,j}\}$, $i = 1, ..., N_\mu$ and $j = 1, ..., N_{\sigma^2}$, if

$$\sup \left\{ \frac{\left| \omega_{i,j} \otimes [D(\mu_i, \sigma_j^2) - \tilde{D}(\mu_i, \sigma_j^2] \right|}{1 + \sup \left| \omega_{i,j} \otimes [D(\mu_i, \sigma_j^2) - \tilde{D}(\mu_i, \sigma_j^2] \right|} \right\} \geq 1.0e^{-3}$$

set $\psi(\mu, \delta, \sigma^2) = \tilde{\psi}(\mu, \delta, \sigma^2)$ and repeat, otherwise stop.

References

Alvarez, F., and R. Shimer (2009a): "Search and Rest Unemployment," *Econometrica*, Forthcoming.

——— (2009b): "Unemployment and Human Capital," Mimeo, University of Chicago.

Barillas, F., and J. Fernández-Villaverde (2007): "A generalizeation of the endogenous grid method," *Journal of Economic Dynamics and Control*, 31, 2698–2712.

Barnichon, R. (2010): "Productivity and unemployment over the business cycle," *Journal of Monetary Economics*, Forthcoming.

Brugemann, B., M. Hagedorn, and I. Manovskii (2010): "Labor Productivity during Postwar Recessions : Basic Data, Facts and Figures," Mimeo.

Burdett, K., S. Shi, and R. Wright (2001): "Pricing and Matching with Frictions," *Journal of Political Economy*, 109(5), 1060–1085.

Carroll, C. (2006): "The method of endogenous gridpoints for solving dynamic stochastic optimization problems," *Economics Letters*, 91, 312–320.

Gali, J., and T. van Rens (2010): "The Vanishing Procyclicality of Labor Productivity," Mimeo.

Gathmann, C., and U. Schonberg (2010): "How General is Human Capital? A Task-Based Approach," *Journal of Labor Economics*, 28(1), 1–49.

Gouge, R., and I. King (1997): "A Competitive Theory of Employment Dynamics," *Review of Economic Studies*, 64(1), 1–22.

Hagedorn, M., and I. Manovskii (2010): "Productivity and the Labor Market : Co-Movement over the Business Cycle," *International Economic Review*, Forthcoming.

Kambourov, G., and I. Manovskii (2008): "Rising Occupational and Industry Mobility in the United States : 1968-97," *International Economic Review*, 49(1), 41–79.

——— (2009a): "Occupational Mobility and Wage Inequality," *Review of Economic Studies*, 76, 731–759.

——— (2009b): "Occupational Specificity of Human Capital," *International Economic Review*, 50(1), 63–115.

Lucas, R. E., and E. C. Prescott (1974): "Equilibrium Search and Unemployment," *Journal of Economic Theory*, 7, 188–209.

Moen, E. (1997): "Competitive Search Equilibrium," *Journal of Political Economy*, 105(2), 385–411.

Moscarini, G., and K. Thomsson (2007): "Occupational and Job Mobility in the US," *Scandinavian Journal of Economics*, 109(4), 807–836.

Peters, M. (2000): "Limits of Exact Equilibria for Capacity Constrained Sellers with Costly Search," *Journal of Economic Theory*, 95, 139–168.

Pissarides, C. A. (2001): *Equilibrium Unemployment Theory*. Cambridge, MA : MIT Press, 2nd edn.

Poletaev, M., and C. Robinson (2008): "Human Capital Specificity : Evidence from the Dictionary of Occupational Titles and Displaced Worker Surveys, 1984-2000," *Journal of Labor Economics*, 26(3), 387–420.

Reiter, M. (2002): "Recursive computation of heterogeneous agent models," Mimeo.

——— (2010): "Solving the incomplete markets model with aggregate uncertainty by backward induction," *Journal of Economic Dynamics and Control*, 34, 28–35.

Robinson, C. (2010): "Occupational Mobility, Occupational Distance and Basic Skills : Evidence from Job-Based Skill Measures," Mimeo.

Shimer, R. (2005): "The Assignment of Workers to Jobs In an Economy with Coordinantion Frictions," *Journal of Political Economy*, 113(5), 996–1025.

Veracierto, M. (2008): "On the cyclical behavior of employment, unemployment and labor force participation," *Journal of Monetary Economics*, 55, 1143–1157.

Young, E. (2010): "Solving the incomplete markets model with aggregate uncertainty using the Krusell-Smith algorithm and non-stochastic simulations," *Journal of Economic Dynamic and Control*, 34, 36–41.