# The task space, task wage gaps and income inequality

J. William Ambrosini Jr*

December 9, 2010

JOB MARKET PAPER

First draft: July 2010[1]

**Abstract**

This paper develops and tests a model of equilibrium task supply where such tasks are the proximate inputs to production. In the model, ex ante identical workers choose task supply and identical, competitive firms choose task demand. The model predicts that changes in task supply adjustment costs drive changes in the income distribution by creating a wedge between task wages. The key assumption of the model — that task adjustment costs are economically significant — is validated using methods of program evaluation from the labor literature. That quasi-experimental evidence indicates workers who are forced to make the median change in their task supply lose on average 7 log points of annual earnings over their lifetimes. Working through the mechanisms of this empirically validated model, increased task adjustment costs can explain much of the recent rise in income inequality.

*I thank my advisers Giovanni Peri, Ann Stevens and Pontus Rendahl for their invaluable feedback and I thank the audiences for talks at UC Davis, Sonoma State University and Sacramento State University and the participants at the All-Cal Labor conference for their questions and comments.

[1]Section 4 largely comes from a working paper for which the first draft was written in early 2009.

The whole of the advantages and disadvantages of the different employments of labour and stock must, in the same neighbourhood, be either perfectly equal or continually tending to equality. If in the same neighbourhood, there was any employment evidently either more or less advantageous than the rest, so many people would crowd into it in the one case, and so many would desert it in the other, that its advantages would soon return to the level of other employments. This at least would be the case in a society where things were left to follow their natural course, where there was perfect liberty, and where every man was perfectly free both to chuse what occupation he thought proper, and to change it as often as he thought proper. Every man's interest would prompt him to seek the advantageous, and to shun the disadvantageous employment.

— Adam Smith, 1776 (1904)

# 1 Introduction: the task space and rising inequality

In this paper, I introduce a non-Ricardian model of the task wage distribution. It is "non-Ricardian", in contrast to the "Ricardian" model of Acemoglu and Autor [2010], because agents are not endowed with different skill levels and so they do not have comparative advantage in supplying high or low skill tasks where such tasks are the proximate inputs into production. Instead, workers in the model choose the bundle of tasks they supply only considering the costs of task adjustment. Task adjustment costs put a wedge between task wages and this feeds through to overall earnings inequality. As these adjustment costs change over time, this leads to changes in the distribution of task wages and this leads, in turn, to changes in the overall distribution of earnings. The purpose of this paper is to show that the non-Ricardian model is a credible story of the increase in overall inequality. To demonstrate its credibility, this paper will lay out and validate the assumptions of the model, derive its optimizing and equilibrium implications and then test that those implications hold in the data.

Task supply adjustments constitute movements in what I call the task space. Mathematically, the task space is a metric space for occupations where the distance between occupations reflects the qualitative difference between occupations in terms of tasks performed on the job. For example, insurance adjusters and economists are near each other in task space while farm product purchasing agents and economists are far from each other in task space. Workers are extremely mobile in the task space in both quantitative and qualitative senses. The average worker makes a substantial change to his bundle of supplied tasks (i.e. he changes "careers") 3.7 times in his working life [2]. Qualitatively, each of these career changes is akin to a vertical move from sales person to marketing manager, a horizontal move from waiter/waitress to kitchen staff or an industry change from farm manager to hotel manager. Another way to quantify these qualitative career changes is look at where in task space workers move. Table 1 shows the transition probabilities for moving into various regions of the task space. These regions were determined by a type of cluster analysis where the five cluster centroids are the five occupations that are furthest from each other in the task space. Moving from one region to the another, then, is an extreme move in the task space. As can be seen, much of the year to year movement in the task space is qualitatively significant. For example, about a fifth of managers change their occupation every year and of those only 11% take jobs as managers and another 11% become "creative" workers (like musicians or kindergarten teachers). In fact, a majority of occupation changes involve transitions between, not within, neighborhoods in the task space. The same story emerges when you analyze transitions between the standard BLS occupation groups (e.g. "Management", "Services", "Sales", etc). This evidence suggests these movements in task space may be economically significant and the non-Ricardian model is a first step in understanding these task supply decisions.

As the only non-standard element of an otherwise standard model, the existence of economically significant task adjustment costs is of prime interest for validating the model. Novel for this type of paper, I use the techniques developed in the applied microeconomics literature over the last 20 years to get credible evidence that these costs are large. In particular, I identify task adjustment costs using displaced workers in the Panel Study of Income Dynamics (PSID) using a generalized difference in difference identification strategy. This analysis shows displaced workers who move the median distance in task space, the treated group, see substantially larger costs of displacement relative to those displaced workers who did not move in task space, the control group. Over their lifetimes, the treated group averages 7 log points lower annual earnings than the control group. Because the

[2] Career changes are defined as more than the median movement in task space. This statistic was calculated using workers from the PSID and assumes the average working life is 45 years.

treatment is not exogenous I use an instrumental variable approach to check for omitted variable bias. The instrument for distance moved in the task space is task isolation before displacement. Analogous to geographic isolation, more task isolated workers have to move further in the task space when they change occupations. The exclusion restriction is that isolation only affects outcomes through task distance, the treatment. I check to see if other direct channels exist, such as task isolation proxying for industry-specific human capital, but only find evidence that the exclusion restriction is valid. The IV estimates using this instrument are not qualitatively different from the OLS estimates which suggests the omitted variable bias is not large.

Using the same techniques, I estimate a time-series for task adjustment costs based on the at-displacement costs of moving in task space. Due to data limitations, this time-series can only be estimated from 1969 to 1993. Over this time, task adjustment costs increased 4.5 fold. However, under the lens of the model, we can infer from changes in the distribution of task wages that task adjustment costs have flattened out since the late 1990s. The other important time-series that will be used in testing the model are task wage gaps and relative task employment. Like adjustment costs, task wage gaps rose substantially since the mid-1970s (see figure 3). Relative task employment (see figure 2) increased somewhat, too, but not nearly as dramatically as costs and relative wages.

The quasi-experimental evidence suggests task adjustment costs are an economically important quantity, but this does not mean that workers interact with them in the ways the non-Ricardian model assumes. Tests that the workers' optimizing conditions hold in the data effectively test that the model is a close approximation to how workers actually choose task supply. Using an error correction model (ECM), I show that the model-implied optimizing condition holds in the long-run. The evidence suggests workers optimally choose their task supply in the way suggested by the model.

Tests of the equilibrium implications of the model, on the other hand, test whether or not changes in the economic environment have the model's predicted effect on task quantities and prices. The first equilibrium implication of the model is that increases in task adjustment costs should decrease the relative supply of tasks. Intuitively, increased adjustment costs make it harder to supply some tasks and easier to supply others. This has the effect of decreasing the supply of the former tasks relative to the latter tasks. This negative relationship is verified in the data in a regression were task wages and a linear trend demand shifter (e.g. task-specific technological change) are among the controls. The control for demand shifts turns out to be crucial because in the raw data relative task supply actually increases slightly. To explain this increase given the large increases in task adjustment costs, there has to be large contemporaneous increases in relative task demand.

The second equilibrium implication is also verified by that regression. Task wage gaps and relative task supply should only be related to each other through task adjustments. In other words, for a given level of adjustment costs there should be zero correlation between relative task wages and relative task employment. In the model, task supply is endogenous so changes in the economic environment unrelated to task adjustment costs will induce changes in task supply. Those task supply changes do not feed back relative task wages as they are completely determined by adjustment costs. In the regression discussed above, adding task wages as a control does not change the coefficient on task adjustment costs and the coefficient on wages is insignificant and reduced when the cost term is added. This result is in contrast to the canonical model of Katz and Murphy [1992] where supply is exogenous and so a significant relationship should exist between wages and supply.

With a credible model — one that has valid assumptions and verified empirical predictions — in hand, the non-Ricardian model can be used to run counter-factual experiments. In particular, holding task adjustment costs fixed and so holding the task wage distribution fixed from 1976 to 2009 but allowing unmodeled factors that affect the income distribution to vary, simulations from the model predict there would have been no increase in the dispersion of income in this period. That is to say: the model can explain all of the increase in inequality experienced in the last few decades.

The rest of the paper is structured as follows. In the next section, I build a non-Ricardian model of task supply choice. In Section 3, I document trends in task employment and then I use those data and wage data to construct trends in task wages. Before I use those quantities and prices to test the optimizing implications of the model in Section 5, in Section 4, I validate the main assumption of the model using techniques of program evaluation[3]. In Section 6, I use the equilibrium conditions of the model to understand the relationship between task employment and task adjustment costs and I find that demand shifters such as "task biased technological change" play a big part in understanding

[3]Section 4 is taken nearly verbatim from Ambrosini [2010a].

trends in task employment. In the subsequent section, Section 7, I show that most of the increase in inequality can be attributed to the mechanisms in the non-Ricardian model. I conclude in Section 8 and I lay out a road map for further research.

# 2 A model of task supply and demand

Rising wage and income inequality over the last three-plus decades is well-documented phenomenon, e.g. Atkinson et al. [2009], Burkhauser et al. [2010] and Heathcote et al. [2010][4]. A number of economists have documented various causes of this rise and they have attempted to quantify the degree to which each cause contributed to it, e.g. Autor et al. [2008] and Card and DiNardo [2002] or Mortensen [2005] and Wong [2003]. Goldin and Katz [2008] summarize and typify most of this literature. They document a race between technology and education (a proxy for presumed fixed and permanent features of workers), with technology winning; the increase in demand for high skills, they say, has outstripped the increase in supply. Their work is typical because it assumes workers are different a priori[5].

The model studied in this section differs in two ways from this literature. First, it treats tasks as the proximate input into production. Second, it assumes agents are ex ante identical.

## 2.1 Preferences

Each worker is endowed with 1 unit of time. A worker supplies her labor inelastically but she chooses how to divide her labor time into supplying the various tasks. She does not prefer supplying any particular task; she will only consider the path of task wages and the costs associated with moving in the task space when determining her task supply mix. She cannot set wages nor does she have a role in determining task switching costs. Her problem is given by,

[4] However, Gordon [2009] provides a number of citations to papers skeptical of the reported magnitudes of the rise. Heathcote et al. [2010] is in a special issue of Review of Economic Dynamics on "Cross Sectional Facts for Macroeconomists".

[5] The model in Mortensen [2005] is predicated on there being a priori identical workers but to match the wage distribution, he has to add heterogeneity to the model (see page 56).

$$\begin{aligned} \max_{\{l_{t+1},h_{t+1},c_t\}_t} \quad & \textstyle\sum_{t=0}^{\infty} \beta^t u(c_t) & \\ \text{s.t.} \quad & l_{t+1} + h_{t+1} & = 1 \\ & \textstyle\sum_{t=0}^{\infty} R^t \left(w_{t,l} l_t + w_{t,h} h_t\right) & = \sum_{t=0}^{\infty} R^t \left(c_t + \hat{D}(h_t, h_{t+1}, l_t, l_{t+1})\right) \end{aligned}$$

She optimizes over the types of tasks ($l$ and $h$) and consumption ($c$). She maximizes her life-time utility, discounting at the rate $\beta$. She has a time constraint and her present value budget constraint is standard except for the cost of changing tasks, i.e. the cost of moving in task space ($\hat{D}$, a function of past and present values of $l$ and $h$). Financial markets are assumed to be complete and $R$ is the assumed constant interest rate, i.e. the economy is small and open. Each worker owns an equal share of the firms but profits will be zero so they are ignored.

Given financial markets are complete, her consumption decision problem can be separated from her task supply problem[6]. Plugging the constraints into the objective function, her problem simplifies to,

$$\max_{h_{t+1}} \sum_{t=0}^{\infty} R^t \left[(w_{t,h} - w_{t,l})h_t - D(h_t, h_{t+1})\right] \tag{1}$$

where $D(h_t, h_{t+1}) = \hat{D}((1-l_t), (1-l_{t+1}), h_t, h_{t+1})$. The following is the necessary condition of optimality,

$$R(w_{t+1,h} - w_{t+1,l}) = \frac{\partial D_t}{\partial h_{t+1}} + R\frac{\partial D_{t+1}}{\partial h_{t+1}}. \tag{2}$$

Clearly, the task premium is increasing in the costs of moving in task space. I will assume the costs of moving in task space are positive and linear in the first difference, i.e. $\frac{\partial D_t}{\partial h_{t+1}} + R\frac{\partial D_{t+1}}{\partial h_{t+1}} = R(\gamma - 1)$[7]. With low-skill tasks the numeraire, the first order condition can be re-written as,

[6] See Acemoglu [2009], page 360. I will ignore the consumption problem in this paper as my only interest is in understanding task supply.

[7] This is not an innocuous assumption. However, there is little evidence to guide me on the form of adjustment costs. Note: the positive adjustment costs assumption makes adjustments costly when workers move away from supplying $l$ tasks. This suggests reducing the relative supply of $h$ tasks comes with negative adjustment costs.

$$\frac{w_h}{w_l} = \gamma \tag{3}$$

or relative tasks are perfectly elastically supplied at time $t$. Because it incorporates the only novel feature of this model, tests to see if this relationship holds in the data are key to validating the model. These tests are performed in Section 5 of the paper.

## 2.2 Production

Constant returns to scale is the only assumption on production that I will utilize in Section 5. This assumption allows me to recover task wages from task employment and wage data. However, to get a better sense for why the wage gap and relative task employment adjust when there is exogenous changes in task adjustment costs, I put further structure on the production side of the economy. This structure will also facilitate the empirical exercises I perform in Section 6.

Labor is an input to an aggregate linear production function, $Y = AL$[8]. I assume aggregate labor is generated from a CES function of aggregate tasks,

$$L = \left[\alpha L_l^{\frac{\sigma-1}{\sigma}} + (1-\alpha) L_h^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}} \tag{4}$$

which, if tasks are paid their marginal product, suggests task wages are,

$$w_h = (1-\alpha)A\left[\alpha\left(\frac{L_h}{L_l}\right)^{-\frac{\sigma-1}{\sigma}} + (1-\alpha)\right]^{\frac{1}{\sigma-1}} \tag{5}$$

and

$$w_l = \alpha A\left[\alpha + (1-\alpha)\left(\frac{L_h}{L_l}\right)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{1}{\sigma-1}}. \tag{6}$$

Taking the ratio of these wage equations gives an expression for relative task demand,

[8]Time subscripts are omitted as the firm will maximize profits period by period.

$$\kappa \equiv \frac{w_h}{w_l} = \frac{1-\alpha}{\alpha}\left(\frac{L_h}{L_l}\right)^{-\frac{1}{\sigma}}. \tag{7}$$

Relative task demand at time $t$ is a downward sloping function of relative task wages at time $t$.

### 2.3 Equilibrium

There is a continuum of workers (indexed by $i$) of measure one. Out of equilibrium, if the wage gap is wide relative to costs, workers will tend to increase their relative supply of $h$ tasks. This would lead to a narrowing of the wage gap per equation (7). With the assumption of linear costs, workers will always be at an interior solution, i.e. they will always adjust their task supply to equalize the wage gap with costs. In particular, there are negative costs from moving out of $h$ tasks and into $l$ tasks. Out of equilibrium, if the wage gap is narrow relative to costs, workers will tend to incur these negative costs (i.e. benefits) as they increase the relative supply of $l$ tasks. In aggregate, this narrows the wage gap. In the equilibrium defined below, however, the distribution of tasks over workers is not determined.

Definition of equilibrium: An equilibrium is a sequence of task prices, $\{w_{t,h}, w_{t,l}\}_{t=0}^{\infty}$ and task employments such that,

1. each worker $i$ chooses task supply, $\{h_{t,i}, l_{t,i}\}_{t=0}^{\infty}$, to solve problem 1,
2. firms choose task demand, $\{H_t, L_t\}_{t=0}^{\infty}$, to maximize profits,
3. and task markets clear, i.e. $L_{t,h} = \int_0^1 l_{t,h,i}di$ and $L_{t,l} = 1 - L_{t,h}$ for $\forall t$.

For the separation theorem to hold, it is assumed the financial and consumption markets are in equilibrium as well. Because neither the workers' FOC (equation (3)) nor equation (7) depend on time, I drop the time subscripts. These two equations and market clearing determine equilibrium in every period.

To derive the general equilibrium comparative statics, I plug the wage equations into the first order condition and simplify to get,

$$\gamma = \frac{1-\alpha}{\alpha}\left(\frac{L_h}{L_l}\right)^{-\frac{1}{\sigma}}. \tag{8}$$

Exogenous increases in the cost of moving in task space decrease relative task supply. To see this, take logs and re-arrange to get,

$$\frac{\partial \log\left(L_h/L_l\right)}{\partial \log \gamma} = -\sigma < 0$$

when $\sigma > 0$. Intuitively, an increase in task adjustment costs makes it harder to move in task space and it discourages transitions from the $l$ tasks to $h$ tasks. This lower relative supply of $h$ tasks drives up the relative wages of those tasks until equilibrium is achieved.

Another equilibrium implication of the model is that because task supply is endogenous and task adjustment costs linear, wages are only associated with supply through task costs. Demand shifts, such as changes in $\alpha$, will not affect wages and for a given level of costs there is no relationship between wages and supply. From the reduced form equation (8) and equation (3), changes in $\alpha$ affect relative supply but not relative wages. For more general forms of task adjustment costs, though, this result will not hold as changes in supply will feed back through costs in the worker's FOC.

Both of these equilibrium conditions are empirically tested in Section 6.

# 3 Task employment: quantities and prices

In this section, I discuss the measurement of task quantities and prices and the trends in those series. Neither task quantities nor task prices are directly observed in the data. For quantities, task employment is constructed by using qualitative occupation data collected by the government to weight hours worked in each occupation. For prices, variation in overall wages and task employment is used to recover task wages.

## 3.1 Measuring task employment levels

O*Net abilities scores (Willison et al. [2008]) are meant to measure "enduring attributes of the individual [in the occupation] that influence performance" which is to say they are meant to describe attributes of occupations that remain stable over time. Each occupation was originally scored by occupational analysts on 52 types of abilities. These abilities are very detailed (e.g. "*The ability to coordinate two or more limbs (for example, two arms, two legs, or one leg and one arm) while sitting, standing, or lying down. It does not involve performing the activities while the whole body is in motion.*").

I use the O*Net ability scores to allocate hours worked in each task. This allocation is best described with an example. Suppose there were 100 hours worked in an occupation that has ability scores *trunk strength* = 32, *inductive reasoning* = 19, *finger dexterity* = 72, etc with the the sum of the ability scores of 2000. In this case, 1.6 hours would be allocated to *trunk strength*, 0.95 hours to *inductive reasoning*, 3.6 hours to *finger dexterity* and so on. Using this method, figure 1 show the trends in task employment in the CPS.

As can be seen, hours worked in all tasks increased over this time period with the overall increase in hours worked. The cyclical nature of overall hours also can be seen in individual task hours. The coloring in these figures reflects the grouping of the tasks into high- (black), medium- (green) and low-skill tasks (red). Tasks were assigned to one of these three groups by examining the tasks performed by the 90th, 50th and 10th earnings percentiles in 1975, respectively. For example, the green tasks were predominantly performed by those in the 50th earnings percentile in 1975. The contrast between trends in overall hours and task hours is best seen in figure 2 where the overall trend has been subtracted out of each time-series.

Here three trends in the task employment data stand out. First, there has been a fanning out of the middle skilled task hours (the green lines). Some of these tasks have declined in employment, relative to trend, while some have increased in employment. Second, there has been an overall decline in low skill task employment (the red lines) and a trend increase in high skill task employment (the black lines). Lastly, the most dramatic changes in task employment mixes (i.e. the order of the lines) happens during recessions. Notably, the trends in high and low skill task employments were partially reversed in the early 2000's recession.

In the analysis in this paper, having 52 time-series to describe task employment becomes unwieldy. Instead, I use the high/medium/low skill task groupings described above. There are other ways to reduce the dimension of this task space. The literature (e.g. Autor and Dorn [2009] and Peri and Sparber [2009]) has tended to somewhat arbitrarily pick and choose subsets of tasks without regard for the overall structure of these data. This tends to introduce systematic biases in these task measures. Using one dimension as a proxy measure while ignoring other dimensions that are highly correlated with it, for example, introduces biased, non-classical measurement error if the other dimensions are correlated with the proxied variable in one direction but correlated with the proxy in the opposite direction. Besides this measurement error bias, this invites data mining as researchers may tend to pick the proxy tasks that maximize in-sample predictive power rather than proxy tasks that are a priori or out-of-sample better proxies. For more details, the interested reader should see Ambrosini [2010b]. Also, in that paper I discuss and implement other methods of dimension reduction, e.g. factor analysis.

## 3.2 Measuring task wages

As with task employment, task wages are not directly observed in the data, but unlike the hours time-series, wages cannot be mechanically constructed. I have to infer task wages by making assumptions about the production structure that connects movements in task employment with movements in overall wages.

Because labor is aggregated with a constant returns to scale function and used in a constant returns to scale overall production function, Euler's Law gives,

$$wL = \sum w_a L_a$$

Dividing this equation by total labor supply, the following is an estimating equation for task wages,

$$w_{i,t} = \sum_{a \in A} w_{a,t} s_{a,i,t} + e_{i,t} \tag{9}$$

where $s_{a,i,t}$ is the share of total hours worker $i$ spent doing task $a$ (from the set $A$)at time $t$. Strictly speaking, Euler's Law holds exactly so there should be no error term and only one observation is needed to identify task wages. The stochastic error term $e_{i,t}$ derives from measurement error (both on the part of the survey respondents, the econometrician and the managers and workers setting wages). Given this error the regression coefficients estimated from regressing wages on task employment shares serve as estimates of task wages.

As can be seen in figure 3, the trend in estimated wages for high-skilled tasks was flat from the beginning of the sample through the mid-nineties. High-skill task wages increased sharply and then by the mid-naughts the trend flattened out again. Low-skill and medium-skill task wages each had a secular decreasing trend through the whole sample period. The net effect was an increase in both the medium-to-high-skill and low-to-high-skill wage gaps and no trend in the low-to-medium-skill wage gap. This approximately corresponds to the trends in 90-50, 90-10 and 50-10 wage gaps, especially towards the end of the sample. The nineties saw an increase in the 90-50 and 90-10 wage gaps, but no trend in the 50-10 wage gap. However, the well-documented pattern in these ratios in the 1980's is not evident in these task wage estimates.

Interestingly, Acemoglu and Autor [2010] report that occupation dummies — as proxies for task employment — explain more variation in wages towards the end of this sample period. This could explain the patterns described in the previous paragraph. If tasks only began to "matter" for explaining wage dispersion starting in the 1990's, we would expect to see the trends in task wages we actually see in that period. Before they mattered, any pattern in task wages would be consistent with the overall trends in wage inequality. In any case, the lack of a trend in task wages in the 1980's suggests that task wages did not go against the overall trends in wage inequality.

# 4 A quasi-experiment: the value of task knowledge[9]

In this section I step back from the structural model and take a closer look at what the data tell us about the costs of moving in task space. To identify these costs and measure their extent, an ideal experiment would randomly select workers and randomly place them in the task space. Treatment in this ideal experiment would be the distance moved in task space. I would want to know if those randomly assigned to move further in the task space see larger costs, in terms of earnings, than those that did not move as far.

The first step in my analysis is to define a task space in which jobs are located. The vector

[9] This section is largely taken from Ambrosini [2010a]

connecting two jobs in this space defines a distance between them and a direction that defines a qualitative difference between them. In general, jobs that are far apart in the space have much different tasks associated with them and jobs that are closer in the space are more similar. I construct this space by using the occupational abilities ratings found in the O*Net database (see Willison et al. [2008]). I then assign distance measures to individuals by matching their occupations to this database.

The second step in the analysis uses the Panel Survey of Income Dynamics (PSID) to show the existence and extent of transferable domain specific human capital. Because this data set is a panel, I can control for individual-specific unobserved factors and I can look at the effect of various treatments (e.g. switching the tasks performed in a job) over a long period of time. This long-run perspective is important because the simple model of transferable human capital[10] is agnostic on the timing of the effects of treatment.

Except as a thought experiment, I cannot run the ideal experiment. The analysis in this section deviates from that ideal experiment in two ways. First, displacements are not random. Lower quality workers are more likely to be displaced. The literature on displacements has dealt with this problem extensively, but the internal validity of my results do not hinge on this issue because in my analysis both the treated and the untreated have been displaced. However, it may be the case that because I am estimating the effects of treatment on a group of workers that have been displaced, my results are not indicative of the effect of treatment on those workers that do not get displaced. Second, placement in the task space post-displacement is not random. To deal with this problem I use a task isolation score (think geographical isolation but in the task space) in an instrumental variable approach.

Two results emerge from my analysis. Job task-specific human capital exists and the more different two jobs are with respect to the tasks performed on the job, the harder it is to transfer that knowledge when workers switch between them. In other words, the costs of moving in task space depend on the distance moved.

[10] See Ambrosini [2010a]

## 4.1 Constructing the task space

The O*Net ability scores define the metric on the task space (see Section 3 above for further discussion of the O*Net abilities and the task space). Occupations are points in this space. The distance between two occupations in this space is called their TASK DISTANCE. This distance is measured using the Mahalanobis distance (Mahalanobis [1936]) which is weighted Euclidean distance where the weights are determined by the correlation structure of O*Net measured task space. This weighting scheme is necessary because many of the O*Net abilities are highly correlated with each other, e.g. occupations requiring high deductive ability often require high inductive ability. These correlations would artificially increase the distance between otherwise similar occupations because the correlated dimensions are in effect being double counted in an unweighted distance measure [11].

The TASK ISOLATION score for an occupation is the average task distance from the occupation to all other occupations in the relevent labor market weighted by the observed supply for each occupation. This is the expected task distance a worker would travel if they were randomly assigned a new occupation in the labor market. Figure 4 shows both the actual distribution of task distance between old and new occupations when workers change occupations and the expected distribution if they just randomly picked new occupations.

TASK TENURE measures the relationship between knowledge of job tasks and the time spent working on those job tasks. For job task $i$, the O*Net database gives a maximum value for knowledge of $i$, $\hat{T}_i^o$, for a worker in a particular occupation, $o$. A worker first observed in the PSID database is assumed to have the average amount of knowledge of $i$,$\bar{T}_i = \sum_o \hat{T}_i^o$. Over time, $T_i$, the amount of tenure in job task $i$ for the worker, approaches $\hat{T}_i^o$as the worker spends time doing that job task. If the worker's current occupation has a maximum knowledge of $i$ greater than $T_i$then $T_i$ is increasing over time. If the worker's current occupation has a maximum knowledge of $i$ less than $T_i$then $T_i$ is decreasing over time. The following is assumed to be the law of motion for $T_i$:

[11] Also, this distance measure, unlike the angular separation for example, is able to distinguish two occupations that require the same proportions of tasks, but differ in magnitudes. For example, suppose both occupation A and occupation B consist of equal measures of analytical tasks and physical tasks, but B requires twice as much of both. Even though both occupations lie on a ray out of the origin of the task space, the Mahalanobis distance would be positive whereas the angular separation would be zero. The latter distance measure systematically underestimates distances between occupations.

$$
\begin{aligned}
T_{i,0} &= \bar{T}_i \\
T_{i,t} &= (1-\alpha)T_{i,t-1} + \alpha \hat{T}_i^{o_t}
\end{aligned}
$$

where $\alpha = 0.25$ puts task tenure at 95% of the maximum after 10 years. This calibration matches the observation that worker's become experts in an occupation after 10 years (as measured by when wages flatten out as a function of occupational tenure). Because job tasks are multidimensional (e.g. 52 dimensions as in the O*Net database), task tenure is multidimensional. A single summary measure of task tenure makes analysis easier. The analysis below uses the maximum task tenure ($\max_i T_i$) as this summary measure. Another obvious choice would have been average tenure ($\frac{1}{52}\sum_i T_i$). The reported results do not depend on this choice.

Workers with task tenure below the median (among workers) are classified as low task tenure. High task tenure workers, symmetrically, are those workers with task tenure above the median.

## 4.2 Displaced workers in the PSID

The Panel Study of Income Dynamics was an annual survey from 1969 to 1995 and has been biannual since. An attempt is made to resurvey the same individuals in each wave. 3.4% of the respondents have been interviewed in every year of the survey, 36.4% of respondents have answered survey questions in at least 10 survey waves and 87.0% of respondents have answered the survey in more than one year. These statistics should not be interpreted to suggest high survey attrition because the survey sample has expanded significantly over the years. As respondents' children grew to be adults with independent households the survey administrators attempted to add the new households to the sample. The sample has also been expanded and contracted over the years with various new initiatives and budget cuts. In 1969, 4,802 heads of household were interviewed and in 2001, 7,574 heads of household were interviewed. In the early 90's, over ten thousand head of households were interviewed.

The survey asks questions about several members of the household, but most data is collected about and via the head of household. The analysis sample is constructed as follows:

- Sample is limited to 1968 to 1999 male heads of household
- Only looking at first displacements
- Removed individuals with reported displacements in 1968 (i.e. those that said their most previous job in the last 10 years was lost because of being displaced)
- Include all individuals that entered the sample
- Exclude records with zero earnings or that don't have task-measurable occupations

Heads of household were used to simplify data scrubbing (non-heads have variables in separate columns). The analysis sample uses only men to avoid the well-known differences between genders in effects of human capital on earnings. However, my results do not seem sensitive to gender as the analysis run on women-only yield qualitatively similar results.

Displacements are defined as narrowly as possible in the PSID. Respondents are asked if they have different employers then they had in the previous survey wave and if so, why. The typical wording of the follow-up question is "what happened with that employer–did the company go out of business, were you laid off, did you quit, or what?". For the purposes of this paper, the respondent is considered having been displaced in the previous year if they answer "company folded/changed hands/moved out of town" or "employer died/went out of business".

Why first displacements? First, it simplifies analysis. Second, there is evidence (see e.g. Stevens [1997]) the effects of subsequent displacements are different than the effects of the first displacement (or workers who get multiply displaced are different from workers that are only displaced once). Previous displacements, but no other factors observable by the econometrician, can predict future displacements. Assuming the same variables are the only factors observable by workers, only the first displacement is unforecastable by the worker. Ignoring the possibility that workers have private information about themselves, but allowing them private information about their firm, we can only assume displacements are exogenous if information about the closure or relocation of the firm has yet to reach the worker. In the analysis in this paper, I assume the first displacement is exogenous two years prior to the displacement. This assumption suggests, for example, that two years before the first displacement, the will-be displaced worker has as much likelihood of being displaced as all other workers. These assumptions about information (particularly the lack of private information) are problematic for the external validity of displacement studies [12], but are common in the literature.

### 4.3 Summary statistics for the analysis sample

Table 2 displays the summary statistics for the analysis sample with the treated and control groups broken out. For nearly every statistic there is no significant difference between the treated and control. However, there is a noticeable difference in earnings with the treated having nearly 10% lower earnings than the control. Interestingly, those workers that end up with higher levels of the treatment (i.e. greater than the median task distance moved post-displacement) have higher earnings than those with lower levels of the treatment. That said, these statistics are calculated over the whole sample period, pre- and post-displacement. This table, then, suggests the treated sample is not that much different from the untreated except for the primary outcome variable.

Table 3 shows some of these statistics two-years before displacement. Here we see the treated are younger than the control at the time of treatment. As such and as a time-varying attribute of the worker, I explicitly control for age in all regressions. That said, the task tenure averages are very similar which suggests the treated have as many years experience working in the same tasks as the untreated. Both groups have similar levels of education as well. Before displacement, as is discussed above the treated have lower earnings than the control. The differences in age cannot account for these differences and it will be assumed in the analysis below that whatever is causing this difference, it is a fixed feature of the individual. Fixed individual effects, then, control for this difference.

Its interesting to note that unlike in table 2, those with greater levels of treatment (i.e. those that move more than the median in task space post-displacement) earn less than those exposed to smaller treatments. This difference accentuates the importance of dealing with the fact that treatment level is endogenous. Workers choosing to move far in the task space seem to have better outcomes in the long run relative to those that choose not to. In Section 4.5 I use task isolation as an instrument for

[12] For example, workers that know they are low quality might self select into occupations and industries that are more likely to experience displacements.

the treatment. Table 3 shows that those with high levels of treatment are more task isolated before displacement than those with low levels of treatment. This shows that while task isolation is plausibly exogenous to outcomes it is correlated with treatment. It is a good candidate to be an instrument for treatment.

### 4.4 The cost to displaced workers of moving in the task space

In this section, the analysis proceeds assuming displacements are exogenous and treatments (occupations or tasks) are exogenous. Changes in occupations and tasks are calculated between those variables' values two years before and after the displacement. This four year interval is long enough before the displacement that earnings have not started to drop and its long enough after the displacement such that most workers have found employment. On the other hand, this window of time is narrow because I want to reduce the number of other factors besides the original displacement (including a second displacement) that may effect changes in outcomes. I want to replicate as much as possible the ideal experiment.

In this experiment, the counter-factual non-displaced group is identified on observations from individuals that were displaced but three or more years before displacement. The control group is displaced workers that did not move in the task space. By in large, these are workers that did not change occupations. The treatment group is displaced workers that changed occupations after displacement. Following Jacobson et al. [1993], Stevens [1997] and Lindo [2009], the following generalized difference in difference model is estimated:

$$w_{i,t} = \sum_s \alpha_s D_{i,s} T_i + \sum_s \delta_s D_{i,s} + X_{i,t}\delta + \beta_i + \gamma_t + \epsilon_{i,t} \tag{10}$$

$w_{i,t}$is the earnings for worker $i$ at time $t$. $D_{i,s}$ is a dummy variable indicating the time of displacement relative to time $t$ and $T_i$ is the treatment of interest. There is one dummy variable for every year from two years before displacement to nine years after. There is a single dummy variable for displacement having happened ten or more years in the past. $X_{i,t}$ is a set of time varying attributes of the worker. A quadratic in age is always among the controls as is this interacted with the treatment measure. The $\beta_i$ is a individual fixed effect and the $\gamma_t$ is a year fixed effect.

Because the most interesting comparison is between the treated and untreated workers (and not between displaced and non-displaced), in the following analysis I will report two statistics. The p-stat is the likelihood, given the imprecision of the estimates, that lifetime earnings of the treated are less than the lifetime earnings of the untreated. The q-stat is the likelihood that earnings in every period post-displacement are less for treated than the untreated. Both statistics are computed using bootstrap simulations.

### 4.4.1 Treatment: occupation switches

The first treatment is occupation switches, a binary treatment. Figure 5 displays the estimates for the effects on earnings of displacement for occupation switchers versus occupation stayers. The red line are occupation stayers and the black line are occupation switchers. Along the x-axis are years since displacement. The reference group, and so the zero line, is the the displaced in the counter-factual case where they were not displaced. This counter-factual is identified using observations of the displaced previous to two years before displacement.

There are a couple potential hypotheses to test. The main prediction of the transferable human capital story regarding this and the following experiments is that given exogenous shifts in the task space, those that move further should lose more human capital and thus see larger wage declines. The theory predicts little about the timing of this loss because it says nothing about when transferable human capital is utilized in a worker's career[13]. This suggests the appropriate hypothesis test is whether the integral of the difference between the red and black lines is positive. The p-stat is the appropriate statistic to test this hypothesis. A much more strict test of this prediction is whether the red line is above the black line in all periods. The q-stat is the appropriate statistic for this test. Evidence for the second hypothesis is necessarily evidence of the first, but as its a much more strict test, it gives a sense of the magnitude of the effect of treatment.

As in Stevens [1997], occupation switchers see much bigger costs from displacements than occu-

[13]Assuming any wage contract is available, a firm might, for example, pay a worker to accumulate human capital through most of his career and then utilize it in a burst of productivity just before the worker retires. The inconsistent productivity of artists and academics is another example.

pation stayers: the p-stat is 95% and the q-stat is 50%. This means its very likely that occupation switchers see a decline in their earnings over their lifetimes after their displacement and there is a very good chance they have lower earnings in every period after displacement.

The evidence for occupation switchers is strong evidence for the existence of transferable human capital. It does not tell us, however, the nature of that sort of human capital. Transferable human capital may be occupation specific, as has been suggested by the previous literature. This is unlikely as the structure of occupation taxonomies (e.g. Census 1990 occ codes) are arbitrarily constructed relative to the actual tasks being done and the products or services being produced on the job. Occupations in these taxonomies are essentially labels on a set of job tasks performed by the people in that occupation. A truck driver for example is someone who drives trucks, has high stamina, is physically fit enough to occasionally lift heavy objects, etc. In some cases the label also identifies the product or, more usually, the service performed by those in the occupation. Doctors are people who have to verbally communicate, use deductive reasoning, etc, but they also provide health care services, for example.

This criticism of the idea that human capital is occupation-specific has two prongs. First, as alluded to above, occupation labels confuse types of transferable knowledge (product knowledge versus knowledge of job tasks). Second, some occupations are more similar to others in terms of the types of tasks performed by people in those occupations. Also, some occupations are similar in the types of products and services being produced. Consider the examples of a truck driver switching occupations to become a taxi cab driver, a nurse becoming a cab driver and a nurse becoming a doctor. Cab drivers do similar things on the job as truck drivers and nurses provide similar services as doctors. A framework that treats transferable human capital as occupation-specific, however, would treat the truck driver becoming a cab driver as equivalent to a nurse becoming a cab driver. In this case, the truck driver doesn't have to learn very many new tasks in her new job that are different from her old job as a truck driver, but the nurse would be required to learn a whole new set of tasks. Also, this framework would treat the nurse becoming a doctor as equivalent to the nurse's switch to being a taxi cab driver. The nurse though is leaving a lot of product domain knowledge behind when he becomes a taxi cab driver that he wouldn't be leaving behind if he were to become a doctor. The occupation-specific human capital framework doesn't distinguish between product knowledge and job task knowledge and it doesn't account for the fact that some occupation moves require more changes in product and task knowledge than other occupation changes.

#### 4.4.2 Treatment: task switch

In response to this criticism of occupation specific human capital, figure 6 displays the estimates for the effects on earnings of displacement for job task switchers (the black line) versus job task stayers (the red line). While treatment in the regression is continuous, in the figure task switchers are somewhat arbitrarily set to be occupation switchers that moved more than the average distance in task space. Changing this threshold, though, doesn't change the results depicted on this figure. While all displaced workers see persistent costs of displacement, task switchers see significantly higher costs of displacement than task stayers (p-stat=100% and q-stat=58%). Another statistic reported in that figure is the mean discounted lifetime difference between the control and treatment. For this treatment, in the bootstrap simulations this is estimated to be 99 log points with a standard deviation of 50.

At first glance you might expect the black line to converge with the red line over time. In the story of transferable human capital, knowledge is acquired through ones career, so you would not expect a one time destruction of human capital (i.e. the displacement) to have permanent effects. Its important to remember, though, that these figures show outcomes relative to the counter-factual non-displaced group. Because age is being controlled for, the counter-factual group consistent of non-displaced workers in ones cohort. Members of ones cohort accumulate human capital at the same pace and so a gap in knowledge persists between the treated and the untreated.

#### Outcomes before displacement

In the ideal experiment the treatment and control groups would not differ in their unobserved characteristics. Unobserved characteristics can not be controlled for and they may have an impact on the treatment effect. One way to infer that the treatment and control groups do not differ in unobservables is to check to see if their outcomes are different before treatment. In figure 7, we see a modification of the analysis on task switches above. As before the red line is the untreated group, the task stayers, and the black line is the treated group, the task movers. Now, however, all of the post-displacement years have been grouped together, we see more years before displacement and the reference group counter-factual is identified on observations of the displaced ten years and more before their displacement. Four and more years before displacement we see the point estimates almost overlap suggesting outcomes in those years were nearly identical. Between three and two years before displacement, the wages relative to the reference group of the untreated appear to be a little higher than those for the treated group but given the error bars not significantly different. This suggests, in terms of outcomes, the treated and control groups were similar before the treatment.

### 4.5 Internal and external validity

The first concern with the experiments in the previous section is one that haunts all studies of the effects of displacements: displacements, even mass lay-offs, are not independent of the quality of workers. Unless the whole plant or firm is shut down, managers have discretion over who to lay-off. Introspection suggests they will lay-off less productive workers. Thus, displacements are not acting causally on outcomes and our estimates of the effect of displacements are biased and probably biased negatively.

While a problem for the displacements literature both in terms of internal and external validity, this is only a problem for this study in regards to external validity. This is because the treatments contemplated in the previous section and the control group where conditional on displacement. The treated and the untreated were both displaced. Assuming the bias is the same for both groups then this will not have an effect my estimates of the affect of treatment.

This issue still brings into question the external validity of my estimates. Its unlikely that workers that are displaced are as a group similar to workers in general in respect to their human capital characteristics. To verify external validity, I ran the experiments in the previous section but limited displacements to those that occurred in counties where unemployment was above 9% (a standard deviation above the national average). These displacements are more likely to be exogenous. The results are qualitatively similar to the results reported in the previous section and increasing the unemployment threshold does not overturn the results either.

However there may be some issues with the internal validity of the estimates in the previous section.

Occupation and task outcomes are not random. Those workers that choose to switch occupations (and thus tasks) are less likely to be harmed by this switch and we do not observe their outcomes if they had not chosen to switch[14].

Using task isolation scores from before the displacement is a way around this problem yet maintain the spirit of the experiments in the previous section. Just as geographically isolated workers will have higher costs in adjusting their labor supply[15], task isolated workers will find more costly to switch industries or occupations. In theory, whether or not a worker was task isolated before the displacement should not have an effect on their earnings after displacement given a particular distance moved in task space. In other words, what should matter for post-displacement earnings is how much task specific knowledge was lost and pre-displacement isolation only affects earnings through this channel by making more likely for isolated workers to move far in the task space post-displacement (conditioned on making a move in task space). This suggests pre-displacement task isolation can be used as an instrument for distance moved post-displacement.

To get a sense for whether or not this exclusion restriction holds, I checked a number of relationships between task isolation and factors known to predict wages. If no relationship is found between these factors and the instrument, this is evidence that the exclusion restriction holds. First, there is extremely low correlation between task isolation and both education and age (-0.00 and 0.01 respectively). A standard deviation increase in task isolation is associated with about a month more education and the same change in task isolation is associated with a month less of potential experience. Workers with the minimum and the maximum values of task isolation observed in the data, all else equal, differ by less than a year in education and experience. Using estimates from the literature, this extreme difference in task isolation would lead to at most a 10% increase in wages.

Because the the exclusion restriction is conditional on individual fixed effects, age and having been displaced, the above results are only suggestive that the exclusion restriction holds. It is best to look for channels through which task isolation affects outcomes in a time-varying way, e.g. industry where

[14]See Ambrosini [2010a] for a discussion of the bias introduced and a derivation of an equation that partitions the estimated effects into the causal effect and bias terms.

[15]Its common in the literature to use labor market geographic isolation (measured by the level or share of employment in a geographical unit) as an instrument. Also, see Blank [2005] for a discussion of the effects of geographic isolation on poverty.

the worker works. Earnings are affected by industry choice and task isolation may affect industry choice in a way unrelated to task distance. For example, task isolation may be a proxy for industry-specific knowledge. If negative industry shocks are driving displacements then displaced workers in task isolated occupations (qua industry-specific human capital) would be more likely to choose to stay in a dying industry (with presumably lower wages and earnings)[16].

However, the within industry variation in isolation scores is much greater than the between industry variation. Sorting industries by mean isolation score, comparing the 95th percentile industry (religious organizations) and the 5th percentile industry (vending machine operators), there is a little over half a standard deviation difference in isolation scores. The ratio of these two statistics and given the $R^2$ in the regression of isolation on industry dummies is 0.18, there seems to be only a minor relationship between industry and task isolation. Task isolation, however, may lead to selection into industries where task isolation has a bigger impact on earnings. Evidence from a saturated regression of earnings on industry dummies and task isolation suggest this is not the case. There is very a low and positive correlation (0.09) between industry mean task isolation and the coefficients on the interaction terms. Notwithstanding the fact that industry is an outcome variable — implying task isolation could be working through task distance to impact industry choice and therefore not violating the exclusion restriction — there appears to be only a moderate amount of interaction between industry and task isolation[17].

In summary, the exclusion restriction seems reasonable. Because the panel estimates control for fixed individual factors, any channel connecting task isolation and outcomes that does not go through task distance would have to be time-varying within the individual. I explored a plausible candidate, industry worked, but I was not able to find evidence that this is the channel through which task isolation has an impact on earnings.

The correlation between task isolation and task distance is 0.40, but there are several endogenous variables on the right hand side of equation (10), one for each year/task distance interaction term

[16]This analysis is complicated by the fact that at-displacement industry is fixed at the individual level and thus is part of the conditioning in the exclusion restriction. Strictly speaking, then, the example does not apply. However, task isolation may interact with industry choice before and after displacement. In Ambrosini [2010a], I show evidence suggesting task-specific human capital is a separate channel from industry-specific human capital.

[17]The industry statistics in this paragraph are weighted by industry employment where industry is at the 3-digit level. Unweighted statistics give nearly identical results.

and for the age/task distance interactions terms. Cleaning the endogenous variables of variation due to the other instruments, the tests of the instruments being zero in the first stage produce chi-square statistics that average 1079 and vary between 43 and 1556. With degrees of freedom of between 13 and 15 on these tests, for the most part the first stage performs well. The only exceptions are the results for $age * tdist$ and $age^2 * tdist$ but the IV bias on these age terms are not of interest. The results from all 15 first stage regressions are reported in the appendix. These large chi-square statistics and the fact that these are just-identified estimates, i.e. they are median-unbiased, suggests the weak instruments critique does not apply in this situation.

Figure 8 reports the second stage estimates. While qualitatively these IV estimates are similar to the OLS results, the error bars are too wide to take away any clear conclusions about the size of the biases of OLS. However, the difference between task stayers and task movers remains significant as these estimates have a p-stat of 96% and a q-stat of 45%.

Task isolation is a relatively successful instrument for task distance. The evidence from this instrument suggests the bias caused by having an endogenous treatment is not large. I have also tried an employment based instrument, a la Neal [1995]. Using his reasoning, if a task has high levels of employment, this crowding will make job search more costly. One definition of task employment is the number of workers the same occupation, but this has a very low correlation with task distance. Another tack is to calculate task employment by taking a weighted average of occupation employment where the weights are an inverse function of task distance. Both of these definitions turn out to be very weak instruments.

### 4.6 Robustness checks and other results

Ambrosini [2010a] contains many robustness checks. The results are robust to various variations in the sample, definitions of the treatment/displacement and outcome variable. There are three robustness checks of note. First, limiting the sample to those workers that have observations in every year does not change the results (and may actually strengthen them) suggesting a limited attrition effect but perhaps a cohort effect. Second, limiting the sample to a 15-year moving window reveals, see figure 9, the effect of treatment has increased over time. Third, there is a strong interaction between pre-displacement task tenure and the effect of treatment. Figure 10 shows those workers with high task tenure see substantial effects of treatment while those with low task tenure, see figure 11, experience almost no effect of treatment. In fact, except for the first year or so after displacement, it is hard to reject the hypothesis that workers with low task tenure see no effect of *displacement*.

Ambrosini [2010a] also discusses whether task-specific human capital is distinguishable from occupation- and industry-specific human capital. The main results reported above implicitly control for occupation knowledge as the entire treated group and many of the control group switched occupations (recall the treatment is not binary like an occupation switch). For the rest of the horse races, the sample sizes are low, but the results suggest task-specific knowledge is important even when controlling for occupation and/or industry knowledge.

### 4.7 Estimation of the task adjustment cost time-series

In the previous section, I showed results suggesting the *lifetime* differential effect of displacement on those that move in the task space changed over time. In this section, I show that the *at-impact* effect changed over time as well. These estimates will be used later in the paper to test implications of the model outlined in Section 2.

To get a sense for the immediate cost of moving in the task space, I estimate equation (10) but only look at the coefficients for the year following the displacement. I obtain the estimate for each year by limiting the sample to the workers displaced in that year. To increase sample size, I also include those that were displaced in a seven year window centered on the instance year. This tends to introduce stationary noise in the time-series that may make some empirical tests more difficult. Figure 12 shows the resulting estimates. The reader should keep in mind that due to sample size issues, these costs are estimated with a tremendous amount of noise. Only a hand full of years have significant estimates different from zero: 1972, 1978, 1981, 1982, 1991 and 1992. That said a clear upward trend is visible with a large dip in the mid-1980s. Also, there seems to be a positive correlation between business cycles and estimated costs (although, sometimes a spike in costs is a leading indicator, e.g. the 1980's recessions, and sometimes these spikes lag the recession dates, e.g. the 1973 recession).

### 4.8 Discussion

Domain specific human capital in the form of task-specific knowledge exists and it is transferable. The analysis in this section has definitively shown that displaced workers that move far in the task space post-displacement relative to their pre-displacement job lose more human capital than those that do not move far in task space. This suggests that it is the task knowledge itself that is valuable.

The size of that value is much less definitive. The uncertainty on the estimates of lifetime cost of moving in the task space, for example, is quite high. Estimates from displaced workers who choose to move the mean distance in the task space after displacement suggest that on an annual basis such a move can cost on the order of 7 log points of earnings. Bootstrap simulations suggest that the error bounds on that estimate, however, are nearly as large as the estimate. In other words, while it is easy to reject the hypothesis that costs are negative, it is not as easy to know precise size of those costs. That said, as was reported in the previous sub-section there is some evidence that the cost of moving in task space has increased since the early 1970s.

Also, there is evidence that task-specific human capital is accumulated. Displaced workers with low task tenure (a measure of the length in time the worker has been doing the same tasks) saw little effect of moving in task space while those with high task tenure saw pronounced effects. This is consistent with a learning-by-doing model of task-specific human capital.

## 5 Testing the optimizing implications of the model

Now that I have shown that the key assumption of the model is valid — there is a cost to moving in task space that is a positive function of the distance moved — below I proceed to empirically validate a key optimizing implication of the model, equation (2). To do this, I will need to make a slight generalization of the model and I will have to make some assumptions about the costs of changing task supply (i.e. the function $D$).

It takes a minor generalization of the model in Section 2 to have the worker choose among $N$ tasks instead of the two assumed in that section. The first order conditions relate task wage *gaps* to costs so there are $\binom{N}{2}$, or $N$ choose two, first order conditions; one for each unique combination of tasks. I will concentrate on only the $N-1$ first order conditions that relate each of the first $N-1$ tasks to the $N$th reference task. These first order conditions are

$$R(w_{t+1,a} - w_{t+1,a'}) = \frac{\partial D_t}{\partial a_{t+1}} + R\frac{\partial D_{t+1}}{\partial a_{t+1}},\ \forall a.$$

This expression is very similar to the equation (2). The only difference is that $a'$ is the reference task and $a$ is one of the $N-1$ other tasks whereas previously $l$ tasks were the reference task and $h$ tasks were the only other type of tasks ($N = 2$ in that case). The left hand side is the discounted task wage gap between tasks $a$ and tasks $a'$. The right hand side is, again, the marginal cost of moving in direction $a$ relative to $a'$. Its important to realize the right hand side can be different for each pair of tasks. This is because each pair represents a different direction moved in task space and, at least potentially, the cost of moving in task space may be different depending on the direction moved.

To get an expression that can be empirically tested, I will make a few assumptions. First, as in the theoretical section, I assume costs are linear in the change in task supply, i.e.

$$D_t = \sum_a \gamma_a \left(a_{t+1} - a_t\right). \tag{11}$$

Under this assumption the right hand side of the FOC becomes $(1-R)(\gamma_a - \gamma_{a'})$ which varies depending on the direction moved relative to the reference task. Second, I assume the gammas are symmetrical in the sense that the right hand sides are proportional to a common cost factor, i.e. there is some constant, $Cost$, such that $(1-R)(\gamma_a - \gamma_{a'}) = \lambda^a * Cost$ for all $a$. With these assumptions about costs, divide each side of the FOCs by $R$ and then take logs to get an expression that looks like,

$$\begin{array}{cccc} gap^a_{t+1} = & \hat{\lambda}^a & + & cost_t \\ & [3.9, 6.1] & & 1.1 \\ & (< 0.1) & & (0.1) \end{array} \tag{12}$$

where $gap^a = \log(w_a - w_{a'})$, $cost = \log Cost$ and $\hat{\lambda}^a = \log \lambda^a$. A strict interpretation of these first order conditions suggests that in every time period (decade, quarter, minute, second), the cost of moving in task space is equal, after multiplying by a constant, to the task wage gap in the next period.

Introspection suggests this will not be true in every instant. Workers do not have perfect foresight, information about changes in costs or the wage gap does not spread instantaneously, adjustment in reaction to new information is not immediate, the variables are not measured perfectly (quite the opposite at least at the level of the econometrician) and so on. These suggest that equation (12) is best thought of as a long term relationship. One way to interpret this that the first order condition is an equality that the economy is always trying to reach, even if it does not attain it.

The starting point for empirical analysis of this long-run relationship is a general dynamic model relating gaps to costs:

$$\begin{array}{llllllllll} gap^a_{t+1} = & \beta^a_0 & + & \beta_1 & cost_t + & \beta_2 & cost_{t-1} + & \alpha & gap^a_t + u^a_t \\ & [1.8, 3.0] & & 1.1 & & -0.8 & & 0.52 & \\ & (< 0.2) & & (0.1) & & (0.1) & & (0.03) & \end{array} \tag{13}$$

where the offset timing of gaps relative to costs is dictated by the theory (i.e. equation 12). This model picks up a large variety of dynamic patterns in the data. A test of the non-Ricardian model will be to see if these dynamics are consistent with the long-run relationship predicted by the model. The data are a panel of 51 task pairs from the O*Net (see Section 3) with the reference task, "Dynamic Strength", chosen to ensure the wage gaps are positive in most years[18]. Task wages come from Section 3 and the common costs come from Section 4.7 (which are measured in log points). The intersection of these two time-series contains 18 years (1976-1993).

It is debatable whether the cost time-series estimated in Section 4.7 is a good measure of the theoretical construct being tested here. These measured costs are the wage costs of moving a median distance in task space for a very particular group of workers (i.e. workers who lost their job because of a plant closure). First, that group may be sufficiently different from most workers that these data are not a good measure of costs in general (i.e. the results in Section 4 are not generalizable). Second, the theory is silent about what exactly are costs. Theory says costs are anything that creates a wedge

[18] A constant is added to all wage gaps to ensure they are positive in all years. The choice of constant does appear to affect the estimates of the long-run equation (equation (12)). The constant was arbitrarily chosen to make the smallest gap close to zero but still positive. The estimates from a non-linear least squares regression of the non-logged long-run equation have estimates (when the log is taken) that are very close to the logged equation with the added constant estimated in the main text. This result — estimates available upon request — suggests the logged equation with a constant is a good approximation to equation (12).

between task wages. These might be direct costs of changing tasks (such as occupational licensing fees or college tuition) or opportunity costs and forgone wages. The cost series I use here is only a measure of forgone wages. Third, in the case of a negative task wage gap, the theory says "costs" are actually gains from moving[19]. The cost series used here are asymmetrical in that sense because they only measure the cost of moving to a less desirable part of the task space. Having noted these issues of using this time-series to test the model, I will forge ahead. This measure of costs is not the whole story, but they do give us a piece of the picture and to test the model I have to take some stand on the measurement of costs.

Now that I have described the data, I will narrate the results of estimation of the long-run relationship and the general dynamic model[20]. The estimated coefficients are shown just below each equation and standard errors are shown in parentheses. In the case that the regression corresponding to equation (12) is not spurious, i.e. that there exists a long-run relationship between gaps and costs, then the estimated coefficients are an estimate of that long-run relationship. The coefficient on costs is not significantly different from one. This is evidence in support of the two assumptions I made about the relationship between measured costs and actual costs, i.e. the data do not reject that costs are linear in the change in supply and that costs are proportional across tasks. The estimates of the gap-specific coefficients of proportionality are all precisely estimated and range between 3.9 and 6.1.

The estimation of the general dynamic model, equation (13), suggests there is a significant relationship between costs and gaps up to one lag of costs (both $\beta_1$ and $\beta_2$ are significant). When more lags are added, those lags are not significant. Looking now at the estimate of $\alpha$, when controlling for costs, the gap is weakly auto-regressive. Adding another lag of the gaps produces a significant coefficient but because of the short time dimension of this panel and because having an additional lag does not affect the other coefficient estimates nor does it affect the analysis below, it is dropped.

If there is a long-term relationship between gaps and costs and both series are integrated[21] (i.e.

[19] For example, a medical doctor becoming a janitor would entail a negative cost.

[20] Estimates reported here are estimated using non-linear least squares using the *nls* function of the R package *stats*. Panel estimates using the *plm* function of the *plm* package give similar results.

[21] An augmented Dickey-Fuller (ADF) test cannot reject that the cost time-series has a unit root even at the 10% level but the Elliott, Rothenberg & Stock (ERS) test rejects a unit root at the 1% level. I will let the graphical evidence adjudicate between this contradictory evidence and I will assume this series has a unit root. ADF tests cannot reject the hypothesis that all of the gap time-series have a unit root and only two of these 51 series have the null rejected using the ERS method. In the latter case, using the method discussed in Hanck 2008, the joint hypothesis that all cost-series

the series are co-integrated), then there exists a error correction representation (ECM) of the more general dynamic model, equation (13). It is convenient to assume this representation has the following form:

$$\begin{array}{lccccc} \Delta gap^a_{t+1} = & \beta^a_0 & + \quad \beta_1 & \Delta cost_t + & \gamma & (cost_{t-1} - gap^a_t) + u^a_t \\ & [1.9, 3.0] & 1.2 & & 0.49 & \\ & (< 0.2) & (0.1) & & (0.03) & \end{array} \tag{14}$$

where $\Delta x_t = x_t - x_{t-1}$ and $\gamma = 1 - \alpha = \beta_1 + \beta_2$ and $\beta^a_0 = \hat{\lambda}^a (1 - \alpha)$ are the two restrictions relative to the general model. The ECM, if true, gives a nice interpretation of the parameters of the general model. Specifically, $\gamma$ is measure of how quickly the economy returns to the long-run equilibrium. Because all variables are in log form, $\gamma = .49$ means the economy returns about half way to equilibrium every year, barring any further disturbances that move the economy away from equilibrium. On the other hand if the economy is in equilibrium, $\beta_1$ tells us how fast the wage gaps react to changes in costs. Because it is greater than one, wage gaps overreact to changes in costs, e.g. a 10% increase in costs would lead to a 12% increase in the wage gap. This suggests that workers adjust their task supply sluggishly, as might be expected.

The ECM is a restricted version of the general dynamic model. Testing if these restrictions hold in the data amounts to a test of whether the theory's predicted long-run relationship holds. The restriction that $1 - \alpha = \beta_1 + \beta_2$ holds in the data with a chi-squared statistic of 2.4 (df=1 and p=12%). Also, the ratio of the $\beta_0$'s to the $\lambda$'s, i.e. $\beta^a_0 / \hat{\lambda}^a$, lie in the relatively tight range of 0.481 and 0.487. With $(1 - \alpha)$ estimated as 0.485 these point estimates suggest the second restriction holds as well. Statistically testing that all these restrictions hold validates these findings (chi-squared statistic of 1.5 with 51 degrees of freedom, p $>$ 0.99)[22].

To summarize, the restrictions of the ECM hold in the data. This suggests the long-run relationship between costs and wage gaps predicted by the model, hold in the data.

have a unit root cannot be rejected.

[22] I assumed the lambdas estimated in equation (12) are fixed and known for this test. Incorporating the uncertainty in those estimates would make it harder to reject these restrictions.

# 6 Increasing adjustment costs and relative task supply

For a particular set of production parameters ($\sigma$ and $\alpha$), the general equilibrium comparative statics in turn imply increases in the adjustment costs will lead to a *decrease* in the relative supply of tasks. The data suggest, on the other hand, that relative task employment *increased* in the period in question. Increases in task adjustment costs, then, must have been accompanied by shifts in relative demand for tasks.

**How large must the demand shifts be?**

In this section I take the increase in adjustment costs as given and then I quantify how large demand shifts would have to be to match observed changes in relative supply. The values of the production parameters are unknown, in fact this paper may be the first time they have been studied. Nevertheless, following Katz and Murphy [1992], I take logs of equation (7) and then assume technological bias (i.e. $\log \frac{1-\alpha}{\alpha}$) can vary as a linear trend over time. As such, I regress log relative task wages on a linear trend and log relative task supply to get:

$$\log(w_h/w_l) = -\underset{(1.45)}{4.32} \log(L_h/L_l) + 0.053\,\text{year} + \text{intercept}, \tag{15}$$

with $R^2 = 0.80$[23]. This point estimate suggests an estimate of sigma of 0.23, i.e. there is strong complementarity between high-skill tasks and low-skill tasks. The estimate on the trend term suggests that technology is becoming more biased towards high-skill tasks at a rate of 5.3% per year. Figure 12 shows the fit of the model. It tends to slightly over-predict relative task wages through 2000 and it does a better job thereafter.

The task adjustment cost series is in consumption cost units. To account for this, I adjust 8 such that low skill task wages are not the numeraire, i.e. I substitute the wage equations (5 and 6) into equation (2) to get,

[23]Task employment and wages are calculated as in Section 3.

$$\gamma = A\left[\alpha\left(\frac{L_h}{L_l}\right)^{-\frac{\sigma-1}{\sigma}} + (1-\alpha)\right]^{\frac{1}{\sigma-1}}\left((1-\alpha) - \alpha\left(\frac{L_h}{L_l}\right)^{\frac{1}{\sigma}}\right) + 1. \tag{16}$$

As a result of this adjustment, the right hand side is a complicated function of relative supply and it includes the Solow residual. From 1976 (the first year I have task employment data) to 1993 (the last year I have cost data), there was an approximate 4.5 times increase in the estimated costs of moving in tasks space (the left-hand side of 16). At the same time the Solow residual increased on the order of 30% suggesting the remaining increase in cost, net TFP, is about 3.5 times. While the point estimates for the intercept and trend terms in the equation suggest a very low alpha, the uncertainty in these coefficients suggest the whole range of that parameter, $\alpha \in [0, 1]$, is plausible given the data. For the moment, I take the ratio of technological bias to be 1.5 in favor of high-skill tasks in 1976. At a rate of increase of 5.3% per year, by 1993 this ratio is 3.6. Just this change in the technology bias explains about 86% of the change in adjustment costs. However, to make up the difference the relative supply would have had to have counter-factually decreased by approximately 5%. In fact, the relative supply of high-skill tasks increased 10% over this period. Thus, my choice of initial technological bias has masked additional demand shifts that resulted in increased relative supply.

With an initial technology bias of 1.2 instead, demand shifters explain too much movement in costs and so the increase in task supply seen in the data brings the change in the right hand side of equation (8) back into alignment with the changes in costs on the the left hand side of that equation. With much larger initial technology biases, ones more in line with task wage share data and the point estimates from equation (15), there has to be unrealistically large reductions (i.e. two orders of magnitude) in the relative task supply to match the data. With smaller initial technology biases ($< 1$), the implied initial costs are counter-factually negative. In summary, this experiment is extremely sensitive to values of alpha, a parameter I know very little about.

Overall, however, the lesson from this exercise is: over a large portion of the parameter space, to explain the small changes in relative task supply found in the data with the somewhat large estimated changes in task adjustment costs, I have to consider dramatic shifts in the demand curve or dramatic changes in what I have been calling the technological task-bias. Even in an economy where task choice is endogenous, task-biased technological change plays an important part in explaining shifts in task supply. In other words, for changes in task adjustment costs to hope to explain wage inequality, these changes must be accompanied by large changes in the task-bias of technology.

**Are task wages independently associated with task employment?**

Equation (8) above shows the equilibrium of the model has implications for the relationship between relative task supply, relative task wages and costs of adjustment. Specifically, labor supply is endogenous and should only be related to wages through adjustment costs. In contrast, in the canonical model of the wage distribution (Katz and Murphy [1992] or Acemoglu and Autor [2010]) relative wage changes and a linear trend in biased technology are directly associated with changes in relative supply. Consider the following equation:

$$\log \frac{L_{t,a'}}{L_{t,a}} = \beta_1 \left( \log \frac{w_{t+1,a'}}{w_{t+1,a}} \right) + \beta_2 cost_t + \mathbf{X_t} \beta_{\mathbf{3}} + \epsilon_{t,(a',a)} \tag{17}$$

This is a regression equation were all the notation has been seen before except $\mathbf{X_t}$, which are demand shifters, and $\epsilon_{t,(a',a)}$, which is the error term specific to the task pair $(a', a)$. Think of this as a general reduced form equilibrium condition where the canonical model predicts relationships between employment and wages in equilibrium (i.e. equation (15)) and the non-Ricardian model predicts relationships between costs and employment (i.e. equation (16)). Estimation of this equation will be used to test predictions both of these models have about various conditional correlations. Both models predict that $\beta_1$ is be negative and both predict that there should be a positive relationship between demand shifters and relative employment. There are a couple predictions the models do not have in common.

**Prediction #1a:** The canonical model predicts $\beta_2$ is zero and the non-Ricardian model predicts $\beta_2 < 0$. These are the comparative statics of each model.

**Prediction #1b:** Under the canonical model, $\beta_2$ should attenuate to zero when wages are added as a control.

**Prediction #2:** Controlling for demand shifters, shifts in supply trace out the demand curve.

This suggests $\beta_1 < 0$. Under the non-Ricardian model supply is shifted by changes in costs and by unmodeled factors. If costs are not controlled for then they are in the error term of the regression. To the extent that costs determine supply, when costs are controlled for, the coefficient $\beta_1$ would attenuate to zero.

Table 6 shows the estimates of various configurations of equation (17). The data used are "high-skilled" and "low-skilled" tasks from Section 3. Wages are also those calculated in Section 3 and costs are the cost time-series from Section 4.7. Following Katz and Murphy [1992] the demand shifters are a linear time trend.

First, for all specifications in the table, the coefficient on costs is negative, as expected in the non-Ricardian model, and in most specifications this estimated coefficient is statistically significant at the 1% level. Comparing specification (3) with specification (1), adding relative task wages to the regression of task supply on costs does not substantively change the estimated coefficient on costs. This suggests, at the very least, costs are an independent channel impacting relative task supply. Second, the coefficient on wages in not significant but comparing specification (2) with (3), adding costs does attenuate the point estimate towards zero as predicted by the non-Ricardian model. In summary, on the two dimensions were the canonical model and the non-Ricardian model give different predictions, the data support the non-Ricardian model.

Some other results: the coefficient on the linear trend demand shifter is always significant but not large. This, again, suggests task biased technological change is important for understanding relative supply shifts or the lack thereof. Also, the coefficient on adjustment costs becomes insignificant when lagged values of the dependent variable are added as in specification (4). The coefficient on the lag term is very large but insignificant. Breusch-Godfrey tests suggest that every specification that includes a term for costs does not have serial correlation in their errors. Specification (2), the only one without the cost term, is the only specification that does not appear to control for this serial correlation. This suggests that serial correlation in the cost term is driving serial correlation in the dependent variable. Specification (5) — stretching the data very thin — adds a lag term for costs. This dramatically decreases the coefficient on the lag relative task supply term relative to specification (4). Also, the estimates on the cost terms in this specification suggest about half of the variation in supply due to costs is coming from memory in the cost series. This may reflect the fact that adjustments to new equilibria takes longer than a period or this may reflect measurement error.

The take away from this section is that task adjustment costs have a significant relationship with task supply while there appears to be little or no relationship between task supply and relative task wages. The predictions of the non-Ricardian model fair better than those of the canonical model. Also, demand shifters continue to be an important part of the story of relative task supply.

# 7 Task wage gaps and rising income inequality

How much can the model tell us about rising overall income inequality? In this section I do a simple counter-factual exercise to see how the rise in overall income inequality can be explained by rising task adjustment costs. Using income regressions like those discussed in Section 3, I calculate a counter-factual income distribution for each year by paying workers their residual from that year plus task wages fixed at their 2009 values. Assuming the model is true, this experiment holds constant the task adjustment costs to the value they had in 2009 and this in turn fixes the task wage distribution and thus the aggregate relative task supply. The only things left to vary are other factors, not in the model, that determine individual incomes.

Given tasks are the proximate inputs to production, there are four factors that determine the overall income distribution over workers:

1. The distribution of overall task productivity (i.e. the number of tasks of any type that can completed per unit of time) over workers

2. The distribution of task supply over workers

3. The distribution of *observed* task wages over tasks

4. The distribution of *unobserved* task wages over tasks

and to the extent the model is true, the distribution of task wages (point 3) is determined by task adjustment costs. Because the model is silent on points 1 and 2 the counter-factual lets the data determine these. Because the data do not let us control for 4, the counter-factual will attribute variation in income coming from task wages to unmodeled factors. On the other hand, to the extent the model is false, the distribution of task wages is not determined by task adjustment costs and the counter-factual is holding fixed unmodeled factors that determine the task wage distribution. Estimating the percentage of variation in the task wage distribution that is explained by task adjustment costs is complicated by the fact that there are 52 choose 2 (= 1326) dependent variables implied by the model's first order condition — one for each task wage gap given the 52 tasks analyzed in this paper. Also, theoretically the relationship between costs and task wage gaps can vary by the direction in task space, i.e. using task adjustment costs of displaced workers as a proxy for costs can mean measurement error is higher for some task gaps and lower for other task gaps[24]. In other words, even if a single statistic could summarize the extent to which task adjustment costs *as measured in Section 4.7* explain variation in the task wage distribution, it would not necessarily explain the extent to which costs, better measured, explain the task wage distribution. Given this caveat, below I assume the model is correct and that holding the wage distribution constant is equivalent to holding task adjustment costs constant.

A common measure of dispersion is the variance of the log of income. This measure is 0.44 and 0.51 in the actual income distributions for 1976 and 2009, respectively. For the counter-factual 1976 income distribution, variance of the log of income is 0.53. In other words, holding the task wage distribution constant explains slightly more than all of the increase in income dispersion from 1976 to 2009.

Figure 15 shows the evolution of the actual and counter-factual variance of log income over time. If all variation in incomes was driven by task adjustment costs through the mechanism in the non-Ricardian model, there would be no variation in the dispersion of income over time if the task adjustment costs did not change. This suggests the counter-factual trend in dispersion should be flat. While the lowess fit on the counter-factual data rejects the relevant hypothesis in a statistical sense, relative to the trend in the actual dispersion the counter-factual dispersion does appear flat. Thus it is hard, in

[24] Regressing each of these 1326 task wage gaps on the cost time series gives $R^2$'s between 0.00 and 0.78. Regressions involving "oral comprehension", "control precision" and "multilimb coordination" have the most variation explained by the cost series and the largest $R^2$ is for the regression of the wage gap between "oral comprehension" and "multilimb coordination" on costs. These regressions are suggestive of what direction in task space the cost time series estimated from displaced workers is best measuring. However, interpreting these $R^2$'s in terms of costs explaining variation in the distribution of task wages is complicated by the fact that because tasks are correlated, costs causing variation on one dimension will cause variation in other directions but these individual regressions will not pick up on this.

an economic sense, to reject the hypothesis, given the model is true, that increasing task adjustment costs are responsible for most, if not all, the increase in overall wage income dispersion.

Another interesting feature of figure 15 is that the counter-factual and actual dispersion series intersect and get relatively flat in the mid- to late-90s. This suggests that the costs of adjusting task supply where increasing up until the mid-90s. This is informative because the series measuring these costs, developed in Section 4.7, ends in 1993. The evidence presented in this section suggests costs of adjustment have not changed much sense 1993.

# 8 Conclusion and further research

In this paper, I evaluated a simple, non-Ricardian model of task supply with ex ante identical workers and tasks as the proximate input to production. The singular novel assumption of the model was that moving in the task space incurs a cost and that cost is increasing in the distance moved. This assumption was validated using a quasi-experiment in the PSID. The continuous treatment of being displaced and moving a distance in the task space was shown to have a significant effect relative to the control of being displaced and not moving in the task space.

When workers take these costs as given and optimize over their bundle of task supply, this implies a particular relationship between task wage gaps and the costs of adjustment. This relationship exists in the data in the long-run. Specifically, the restrictions implied by a error correction model on a general dynamic model hold in the data. The estimates of the error correction model suggest the economy moves toward equalizing the first order conditions of the model relatively rapidly. While this finding of a long-run relationship between task wage gaps and the costs of adjustment validates the optimizing framework assumed in the model, the estimates also imply workers move somewhat sluggishly in updating their bundle of supplied tasks. This, in turn, implies there is a very short run (on the scale of few months) in which models with exogenous task supply may better explain the data. Otherwise, at the scale of a year or longer, the non-Ricardian model explains the data.

When the non-Ricardian model has different equilibrium predictions than the canonical model outlined in Katz and Murphy [1992] and Acemoglu and Autor [2010], the non-Ricardian model's predictions are born out in the data. Task adjustment costs have an economically and statistically significant negative relationship with relative task employment in the data. As reported in Section 6, this relationship does not disappear nor is it diminished when controls for relative task wages are added. These successes of the non-Ricardian model relative to the canonical model suggests distinguishing between skills (e.g. education) and tasks is empirically important, as is endogenizing task supply.

While the underlying mechanisms are dramatically different, the non-Ricardian model and the canonical model do share an important general equilibrium implication. Both models imply large demand shifts, not just supply shifts, are required to explain observed prices and quantities. Specifically, for the non-Ricardian model to be consistent with the task employment data and to explain trends in the task wage distribution, increases in task adjustment costs must be accompanied by significant increases in the task bias of technology. As a mere coincidence, these two trends are unsatisfying as an explanation of rising income inequality. I leave it to future research, then, to find mechanisms that link increasing adjustment costs to the task-bias of technological change.

Finally, the simple counter-factual exercise in Section 7 suggests changes in the task wage distribution can explain most of the increase in income inequality since 1976. Through the lens of the non-Ricardian model, it is increased task adjustment costs that are ultimately responsible for this increased inequality.

# References


D. Acemoglu. *Introduction to modern economic growth.* Princeton Univ Press, 2009.

D. Acemoglu and D. Autor. Skills, Tasks and Technologies: Implications for Employment and Earnings. *Handbook of labor economics, forthcoming*, 4, 2010.

J. Ambrosini. Does task-specific human capital exist and is it transferable? Evidence from displaced workers in the PSID. *WORKING PAPER*, 2010a.

J. Ambrosini. Long-run trends in task employment and wages. *WORK IN PROGRESS*, 2010b.

A.B. Atkinson, T. Piketty, E. Saez, and National Bureau of Economic Research. *Top incomes in the long run of history.* National Bureau of Economic Research Cambridge, Mass., USA, 2009.

D. Autor and D. Dorn. Inequality and Specialization: The Growth of Low-Skill Service Jobs in the United States, 2009.

D.H. Autor, L.F. Katz, and M.S. Kearney. Trends in US wage inequality: revising the revisionists. *The Review of Economics and Statistics*, 90(2):300–323, 2008.

R.M. Blank. Poverty, policy, and place: How poverty and policies to alleviate poverty are shaped by local characteristics. *International Regional Science Review*, 28(4):441, 2005.

R.V. Burkhauser, S. Feng, S.P. Jenkins, and J. Larrimore. Estimating trends in US income inequality using the current population survey: the importance of controlling for censoring. *Journal of Economic Inequality*, pages 1–23, 2010.

D. Card and J.E. DiNardo. Skill-biased technological change and rising wage inequality: Some problems and puzzles. *Journal of Labor Economics*, 20(4):733–783, 2002.

C.D. Goldin and L.F. Katz. *The race between education and technology.* Belknap Press, 2008.

R.J. Gordon. Misperceptions About the Magnitude and Timing of Changes in American Income Inequality. *NBER Working Paper*, 2009.

J. Heathcote, F. Perri, and G.L. Violante. Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967-2006. *Review of Economic Dynamics*, 13(1):15–51, 2010.

L.S. Jacobson, R.J. LaLonde, and D.G. Sullivan. Earnings losses of displaced workers. *The American Economic Review*, pages 685–709, 1993.

L.F. Katz and K.M. Murphy. Changes in relative wages, 1963-1987: Supply and demand factors. *The Quarterly Journal of Economics*, 107(1):35–78, 1992.

Jason Lindo. Are Children Really Inferior Goods? Evidence from Displacement-driven Income Shocks. *Journal of Human Resources, forthcoming*, 2009.

P.C. Mahalanobis. On the generalized distance in statistics. In *Proceedings of the National Institute of Science, Calcutta*, volume 12, page 49, 1936.

D.T. Mortensen. *Wage dispersion: why are similar workers paid differently?* The MIT Press, 2005.

D. Neal. Industry-specific human capital: Evidence from displaced workers. *Journal of labor Economics*, pages 653–677, 1995.

G. Peri and C. Sparber. Task Specialization, Immigration and Wages. *American Economic Journal: Applied Economics*, 1(3):135–169, 2009.

A. Smith. An Inquiry into the Nature and Causes of the Wealth of Nations. *Library of Economics and Liberty*, 1776 (1904).

A.H. Stevens. Persistent effects of job displacement: The importance of multiple job losses. *Journal of Labor Economics*, pages 165–188, 1997.

S. Willison, S. Tsacoumis, and C. Byrum. *ONET analyst occupational abilities ratings: analysis cycle 7 results.* Human Resources Research Organization, 2008.

L.Y. Wong. Can the Mortensen-Pissarides model with productivity changes explain US wage inequality? *Journal of Labor Economics*, 21(1):70–105, 2003.

Table 1: Task space yearly conditional transition probabilities

| | Management | Creative | Organizing | Support | Manual labor | n |
|---|---|---|---|---|---|---|
| Management (21%) | 11% | 11% | 26% | 31% | 21% | 2565 |
| Creative (18%) | 14 | 7 | 19 | 37 | 23 | 2325 |
| Organizing(20%) | 16 | 10 | 17 | 30 | 27 | 4268 |
| Support (15%) | 13 | 15 | 23 | 25 | 24 | 5928 |
| Manual labor (20%) | 11 | 10 | 24 | 29 | 26 | 4676 |
| Total | 13 % | 11 % | 22 % | 29 % | 25 % | 19762 |

NOTE: Rows are the regions in task space before workers changed occupations and columns are the regions workers went to after they changed occupations. The percentage in the cell is the percent of occupation switches that are from row occupations to column occupations, e.g. 14% of creative workers that changed occupations became managers. "Management" includes occupations such as CEOs, real estate managers and civil engineers. "Creative" includes marketing managers, musicians and kindergarten teachers. "Organizing" includes accountants, data entry keyers and curators. "Support" includes service worker supervisors, office machine operators and inventory clerks. "Manual labor" includes police men and women, farm workers and machine operators. Data come from the PSID between 1968 and 1999: male heads of household with positive earnings that changed occupations. Due to significant measurement error in occupation transitions post-1980, the unconditional occupation switch rate, shown in parenthesis in the first column, are estimated using data from before 1981. The conditional transition rates reported in the cells are not sensitive to this problem.

Table 3: Summary Statistics two years before displacement

| Data | Control | Treated (< median) | Treated (> median) |
|---|---|---|---|
| age | mean: 37.5 | 34.7 | 34.2 |
| | (SD: 11.4) | (10.4) | (10.8) |
| years of education | 11.4 | 11.6 | 11.5 |
| | (2.7) | (2.8) | (2.4) |
| Occupation isolation | 41.5 | 34.0 | 51.4 |
| | (16.4) | (8.6) | (21.0) |
| Task tenure | 0.77 | 0.77 | 0.78 |
| | (0.13) | (0.12) | (0.12) |
| Hours worked | 2160 | 2183 | 2158 |
| | (710) | (689) | (756) |
| Earnings (82-84$) | 19,392 | 18,156 | 17,716 |
| | (16,154) | (14,300) | (14,318) |

Note: Mean values are reported for task stayers (the control group), task changers were the task distance was less than the median and task changers where the task distance was greater than the median. Standard deviation is reported in parenthesis.

Table 2: Summary Statistics for Analysis Sample

| Data | Whole sample | Non-displaced | Control | Treated |
|---|---|---|---|---|
| age $\in [17, 99]$ | mean: 39.0<br>(SD: 12.4) | 39.2<br>(12.7) | 38.2<br>(11.4) | 38.7<br>(11.4) |
| gender | male<br>(0) | male<br>(0) | male<br>(0) | male<br>(0) |
| years of education $\in [0, 19]$ | 11.8<br>(2.9) | 11.8<br>(3.0) | 11.7<br>(2.7) | 11.6<br>(2.6) |
| number of observations $\in [1, 34]$ | 20.5<br>(9.7) | 19.5<br>(9.9) | 22.8<br>(8.8) | 25.2<br>(7.6) |
| Occupation isolation$\in (16, 131)$ | 44.0<br>(18.2) | 44.3<br>(18.5) | 42.5<br>(17.5) | 42.5<br>(17.2) |
| Task tenure$\in [0, 1]$ | 0.78<br>(0.14) | 0.78<br>(0.14) | 0.77<br>(0.15) | 0.80<br>(0.12) |
| Ever displaced | 23.4% | 0% | 100% | 100% |
| Displaced year prior to survey | 1.2% | 0% | 5.7% | 4.8% |
| Switch 3-digit occupation prior year \| displaced year prior | 65.3% | - | 0% | 100% |
| Task distance \| displaced year prior, occ switch$\in [0, 204)$ | - | - | - | 78.6<br>(29.6) |
| Hours worked $\in [0, 7800]$ | 2126<br>(709) | 2119<br>(705) | 2135<br>(737) | 2,118<br>(698) |
| Earnings (82-84$)$\in (0, 653300]$ | 20,301<br>(17,960) | 20,584<br>(17,955) | 19,715<br>(18,752) | 18,169<br>(17,853) |
| Hours \| task distance greater than median post displacement | - | - | - | 2164<br>(697) |
| Earnings \| task distance greater than median post displacement | - | - | - | 18,880<br>(15,873) |

Note: Mean values are reported for the whole sample, those workers who were never displaced, task stayers (the control group) and those that moved the median distance in task space (the treated group). Standard deviation is reported in parenthesis.

Table 4: Unit root tests: costs for displaced workers

| | $Gap$ | $\hat{Gap}$ | $Cost$ | $\hat{Cost}$ | $z$ |
|---|---|---|---|---|---|
| DF | X | X | X | X | X |
| ADF (k=1) | X | X | X | X | X |
| ADF (k=2) | X | X | X | X | X |
| ADF (k=1, drift) | X | X | X | X | ***X*** |
| ADF (k=1, trend, drift) | X | ***X*** | X | ***X*** | X |

NOTE: Each cell contains the results of a variation on the Dickey-Fuller test for unit root. The hat is the log of the variable and z is the cointegration relationship being tested. The null hypothesis is that the series has a unit root. X indicates the null cannot be rejected, one star is that the null can be rejected at the 10% level, two stars 5% and three stars 1%. DF means Dickey-Fuller, ADF means augmented Dickey-Fuller (with k lagged difference terms), drift means a constant term is added and trend means a trend term is added. The underlined and bolded cells are the author's a priori preferred tests.

Table 6: Task supply on task wages and adjustment costs

| Dependent variable: log rel. task supply | (1) | (2) | (3) | (4) | (5) |
|---|---|---|---|---|---|
| 100 * Costs (units: log earnings) | -5.5 *** | - | -5.3 *** | -3.2 | -2.8 |
| | (1.6) | | (1.7) | (2.7) | (2.5) |
| 100 * Costs (t - 1) | - | - | - | - | -2.3 |
| | | | | | (2.5) |
| 100 * log rel. task wage | - | -3.1 | -2.2 | -2.3 | -2.5 |
| | | (2.7) | (2.1) | (2.4) | (2.4) |
| 100 * year | 0.8 *** | 0.7 *** | 0.8 *** | 0.5 ** | 0.6 ** |
| | (0.04) | (0.04) | (0.04) | (0.2) | (0.2) |
| 100 * log rel. task supply (t - 1) | - | - | - | 32 | 6.9 |
| | | | | (27) | (11) |
| Intercept | -14 *** | -14 *** | -14 *** | -9.9 ** | -12 ** |
| | (0.7) | (0.9) | (0.7) | (4.0) | (4.5) |
| N | 18 | 18 | 18 | 17 | 17 |

NOTE: Employment and wage data are described in Section 3 and the cost data are described in Section 4.7. Data are from 1976 (the first year for employment/wage data) to 1993 (the last year for cost data). All specifications, except (2), produce residuals without serial correlation (the null of no serial correlation was not rejected in the Breusch-Godfrey test).
*** - significance at the 1% level, ** - significance at the 5% level, * - significance at the 10% level

Table 5: Relationship between high and low task wage gaps (rel. to medium tasks)

| Dependent variable: $Gap^h$ | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) |
|---|---|---|---|---|---|---|---|---|---|
| $Gap^l$ | -3.91 | -3.99 | .83 | .57 | 1.56 *** | 2.80 | 1.67 ** | .33 ** | 1.14 *** |
| | 4.29 | 3.83 | .61 | 1.15 | .71 | 3.35 | .60 | .15 | .28 |
| trend | - | .03 * | .59 *** | -.002 | .48 *** | -.02 | -.01 ** | - | .32 *** |
| | | .02 | .08 | .008 | .15 | .02 | .005 | | .10 |
| $Gap^h$1 year lag | - | - | - | .95 *** | .36 * | .64 | .48 ** | .55 *** | .32 ** |
| | | | | .21 | .20 | .44 | .23 | .18 | .15 |
| $Gap^h$2 year lag | - | - | - | - | - | .77 | .10 | .09 | -.02 |
| | | | | | | .79 | .26 | .20 | .16 |
| $Gap^h$3 year lag | - | - | - | - | - | - | .68 ** | .26 | .32 * |
| | | | | | | | .09 | .17 | .17 |
| intercept | 62 | - | -1163 *** | - | -962 *** | - | - | - | -633 *** |
| | 41 | | 156 | | 294 | | | | 200 |
| Years | 24 | 24 | 24 | 23 | 23 | 23 | 22 | 22 | 22 |
| Durbin-Watson stat | 0.3 | 0.3 | 0.7 | 1.9 | 1.4 | 1.5 | 1.1 | 1.2 | 1.3 |

NOTE: $Gap^l$ is instrumented using lagged values, minimum wage and the proportion of new immigrant workers to total workers. First-stage is reported in the text. Each column represents an alternative way to control for non-stationarity in the dependent variable.

*** - significance at the 1% level, ** - significance at the 5% level, * - significance at the 10% level

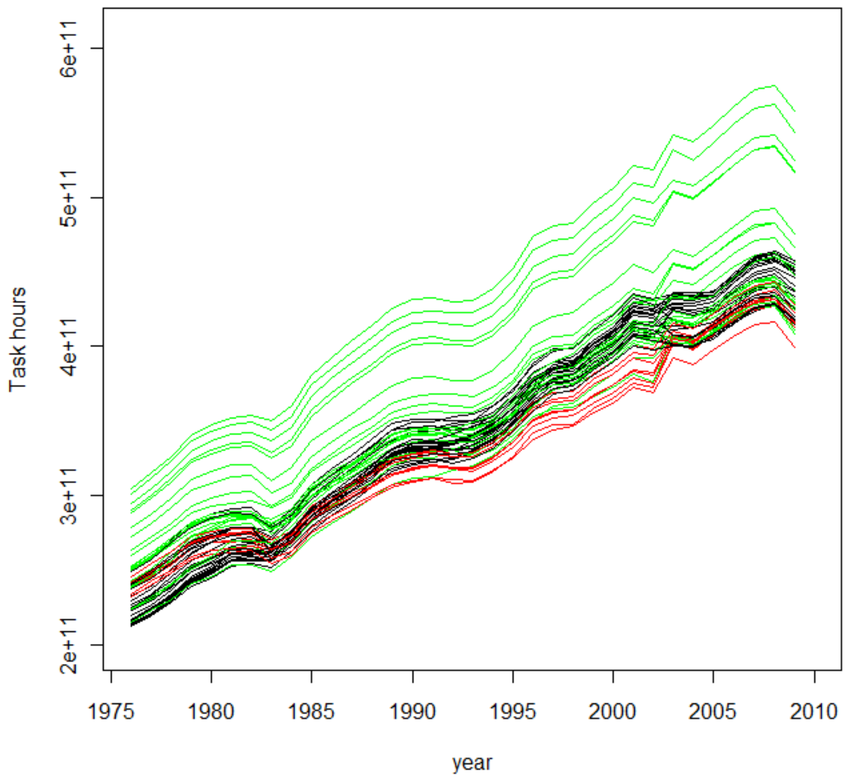


Figure 1: Disaggregate task employment in hours for all 52 O*Net abilities
NOTE: The lines have been colored to show the trends in three groups of tasks. Red are tasks that were primarily performed by those in the 10th percentile of earnings in 1976. Black are the tasks primarily performed by those in the 90th percentile and the green are tasks primarily performed by those in the 50th percentile.

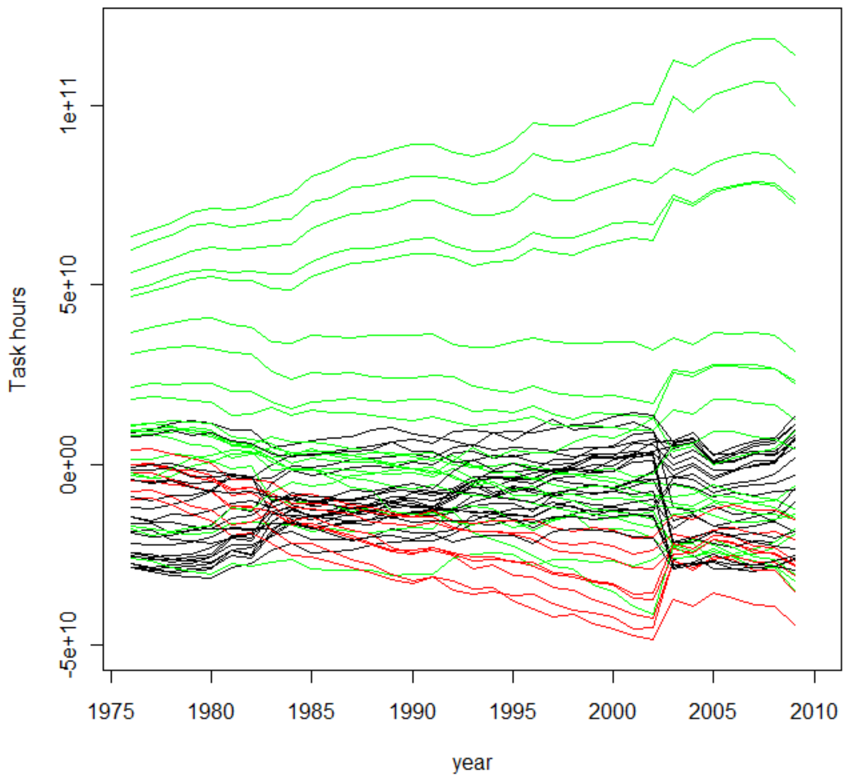


Figure 2: Detrended disaggregate task employment in hours for all 52 O*Net abilities
NOTE: The lines have been colored to show the trends in three groups of tasks. Red are tasks that were primarily performed by those in the 10th percentile of earnings in 1976. Black are the tasks primarily performed by those in the 90th percentile and the green are tasks primarily performed by those in the 50th percentile.

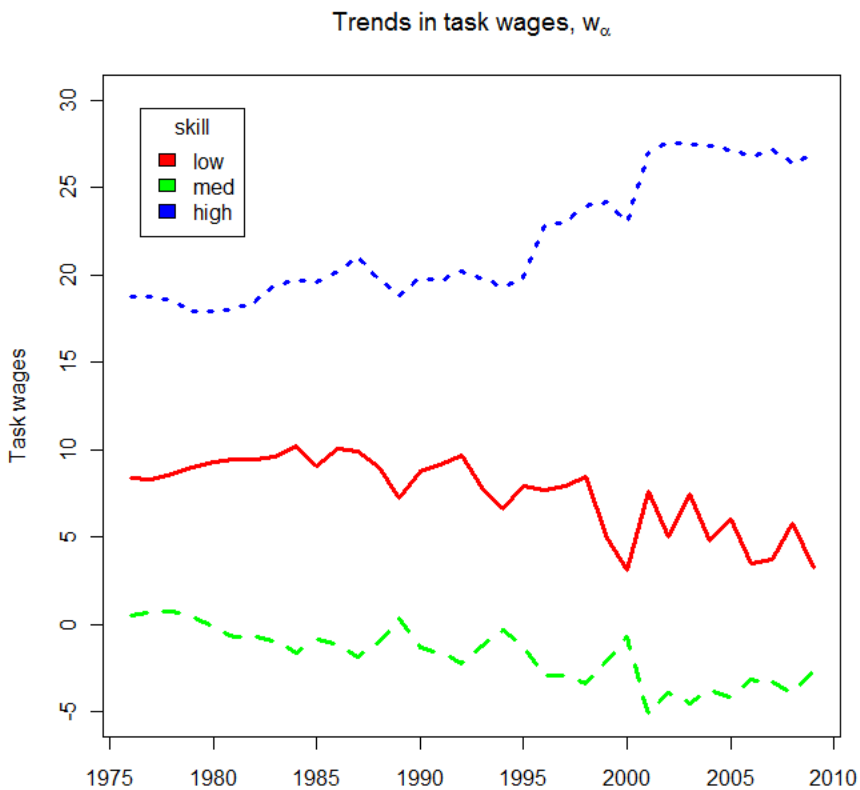


Figure 3: Task wages: $\alpha$

NOTE: The red line is the estimated task wages for tasks that were primarily performed by those in the 10th percentile of earnings in 1976. Blue is the task wages for the tasks primarily performed by those in the 90th percentile and the green line is task wages for those tasks primarily performed by those in the 50th percentile. Wages were estimated for each year using equation (9).

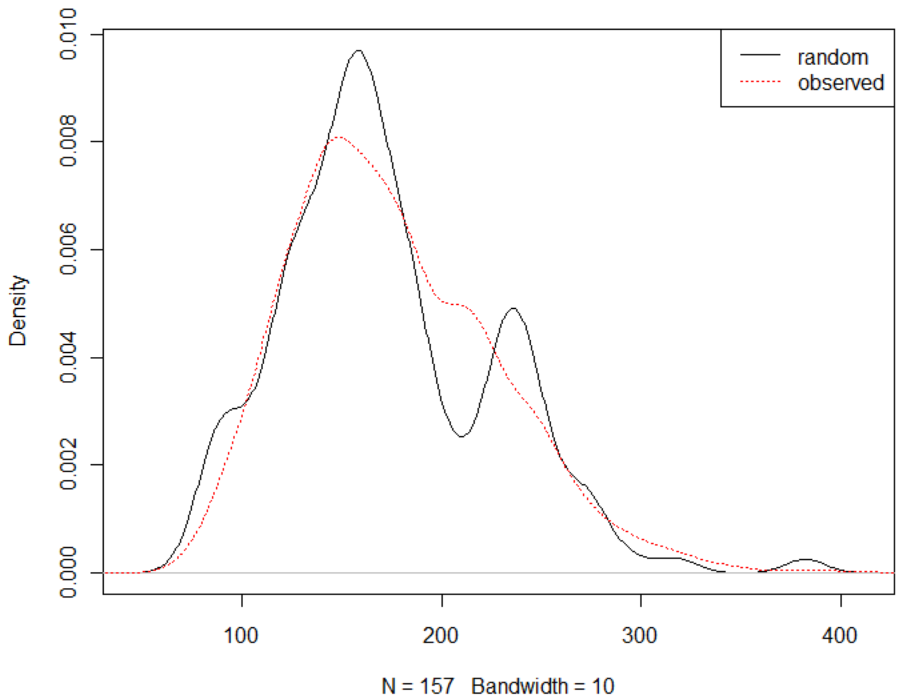


Figure 4: The density of the task distance between old and new occupations for occupation switchers. The black solid line is the expected density if new occupations were assigned randomly. The red dotted line is the actual density of occupation movers.

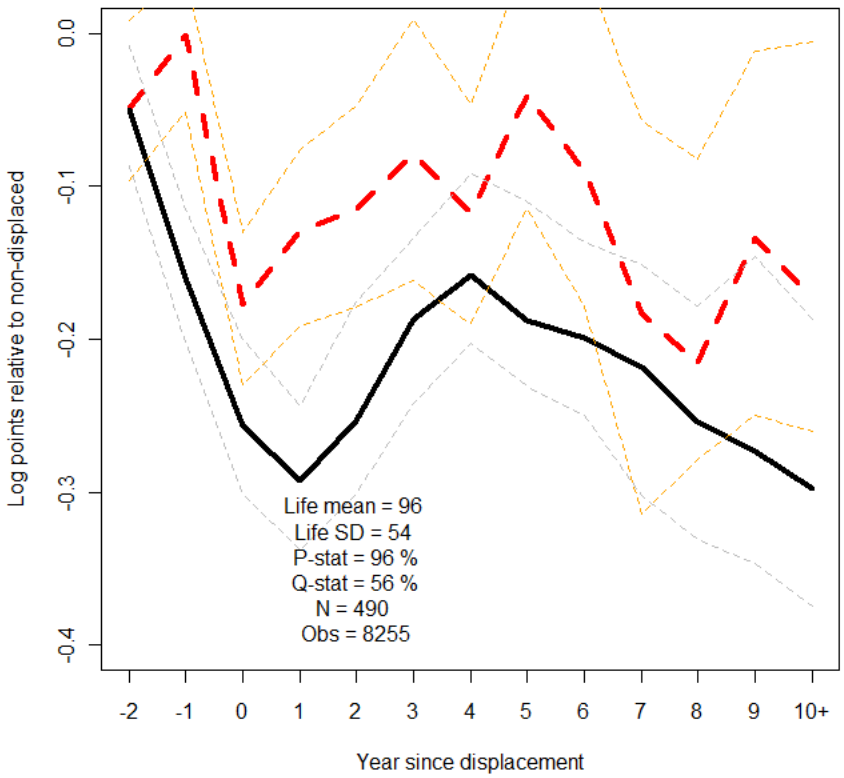


Figure 5: Effects on earnings of displacement for occupation switchers (vs. stayers): Coefficients on displacement dummies from estimating equation (10) on the analysis sample. The analysis sample is described in Section 4.2. An age quadratic is the only control. The thick red (dotted) line represents the estimates on the occupation stayers (three digit occupation two years after the displacement is the same as two years before displacement) and the thick black line represents the estimates on the occupation switchers. Lines representing the confidence intervals (±2∗S.E.) of these estimates are also plotted. Standard errors are panel robust.

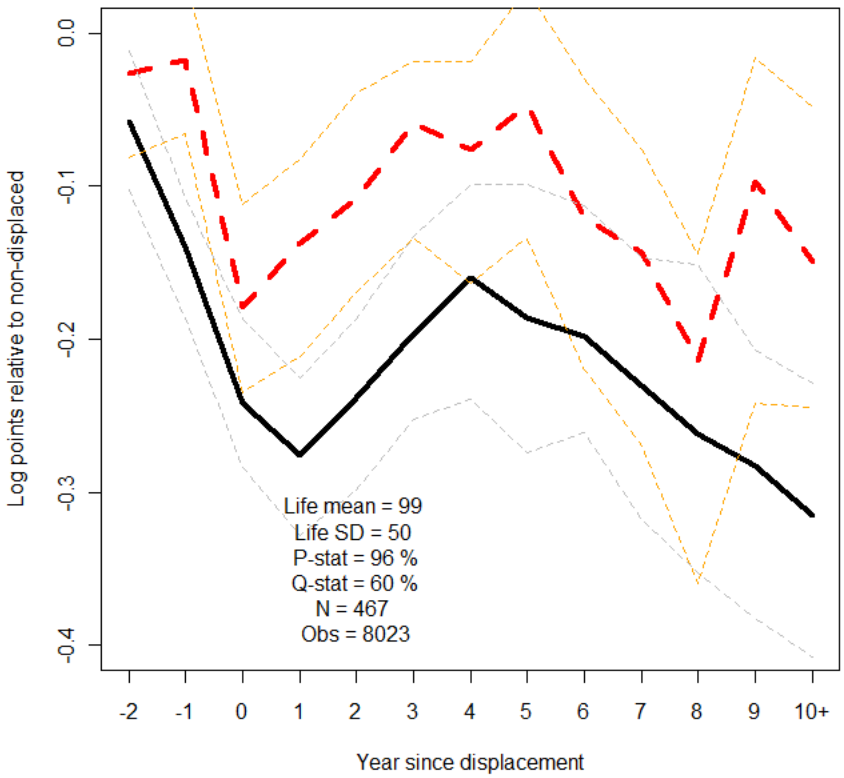


Figure 6: Effects of displacement on job task switchers (vs. non-switchers): Coefficients on displacement dummies from estimating equation (10) (stratification on task distance moved) with labor earnings as the dependent variable. This regression used the analysis sample which is described in Section 4.2. An age quadratic is the only control. The thick red (dotted) line is the marginal effects for non-task switchers and the thick black line is the marginal effects for task switchers who moved the mean distance in task space. Lines representing the confidence intervals (±2∗S.E.) of these estimates are also plotted. Standard errors are panel robust.

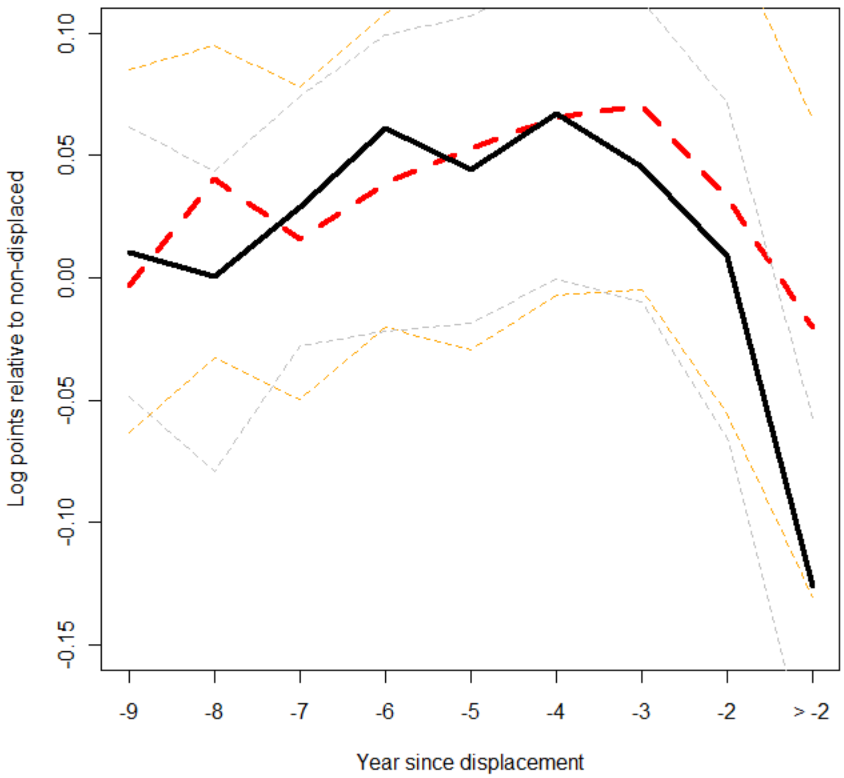


Figure 7: Pre-displacement outcomes for job task switchers (vs. non-switchers): Coefficients on displacement dummies from estimating equation (10) (stratification on task distance moved) with labor earnings as the dependent variable. This regression used the analysis sample which is described in Section 4.2. An age quadratic is the only control. The thick red (dotted) line is the marginal effects for non-task switchers and the thick black line is the marginal effects for task switchers who moved the mean distance in task space. Lines representing the confidence intervals (±2∗S.E.) of these estimates are also plotted. Standard errors are panel robust.

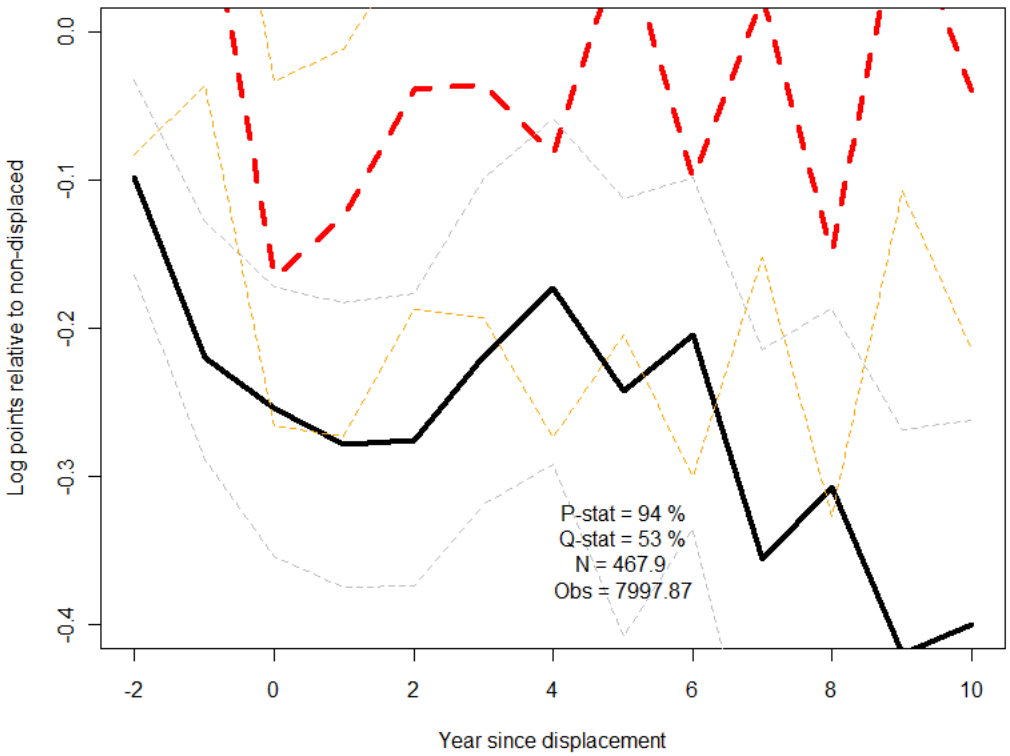


Figure 8: Effects of displacement on job task switchers (vs. non-switchers) with pre-displacement task isolation as an instrument for task distance: Coefficients on displacement dummies from estimating a stratified version of equation (10) (stratification on task distance moved) with labor earnings as the dependent variable. This regression used the analysis sample which is described in Section 4.2. An age quadratic is the only control. The thick red (dotted) line is the marginal effects for non-task switchers and the thick black line is the marginal effects for task switchers who moved the mean distance in task space. Lines representing the confidence intervals (±2∗S.E.) of these estimates are also plotted. Standard errors are panel robust.

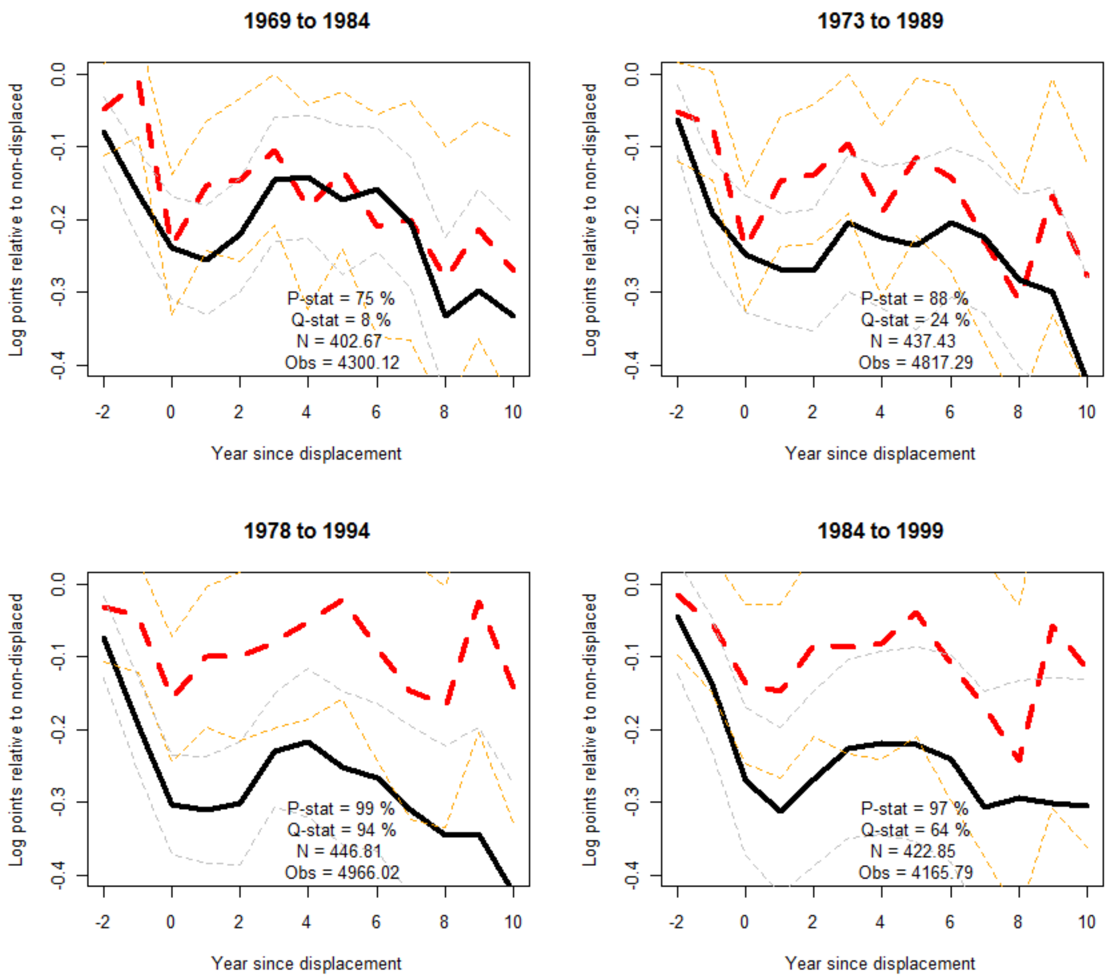


Figure 9: Time restricted samples; effects of displacement on job task switchers (vs. non-switchers): Coefficients on displacement dummies from estimating a stratified version of equation (10) (stratification on task distance moved) with labor earnings as the dependent variable. These regressions are on subsets of the analysis sample — which is described in Section 4.2 — which consist of a moving window of 15 years. An age quadratic is the only control. The thick red (dotted) line is the marginal effects for non-task switchers and the thick black line is the marginal effects for task switchers who moved the mean distance in task space. Lines representing the confidence intervals ($\pm 2*$S.E.) of these estimates are also plotted. Standard errors are panel robust.

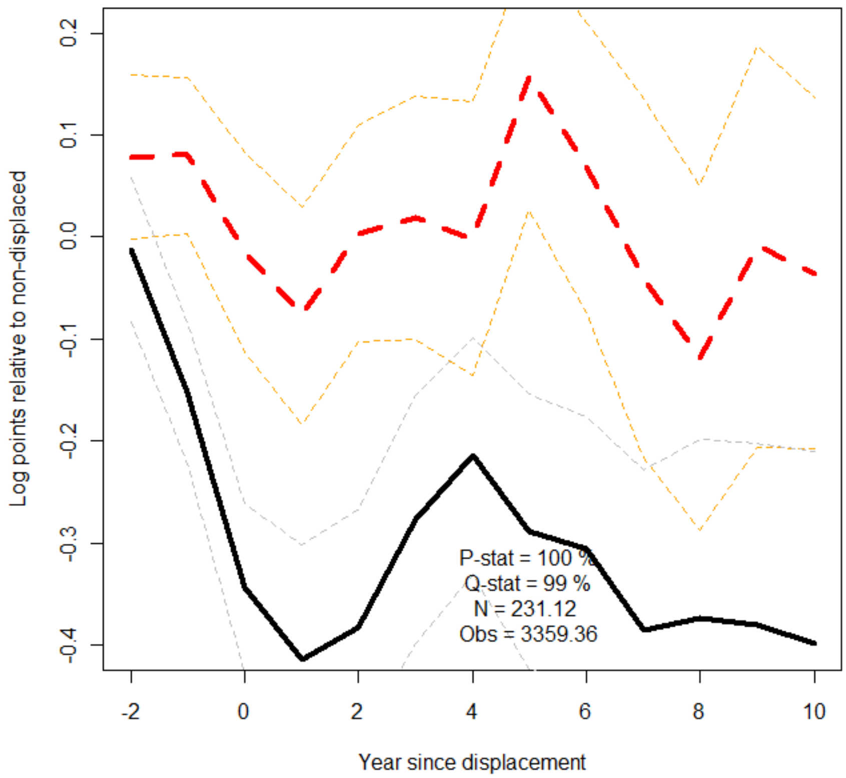


Figure 10: High task tenure sample; effects of displacement on job task switchers (vs. non-switchers): Coefficients on displacement dummies from estimating a stratified version of equation (10) (stratification on task distance moved) with labor earnings as the dependent variable. This regression used a subset of the analysis sample — which is described in Section 4.2 — which consists of individuals that had high task tenure before displacement. An age quadratic is the only control. The thick red (dotted) line is the marginal effects for non-task switchers and the thick black line is the marginal effects for task switchers who moved the mean distance in task space. Lines representing the confidence intervals ($\pm 2*$S.E.) of these estimates are also plotted. Standard errors are panel robust.

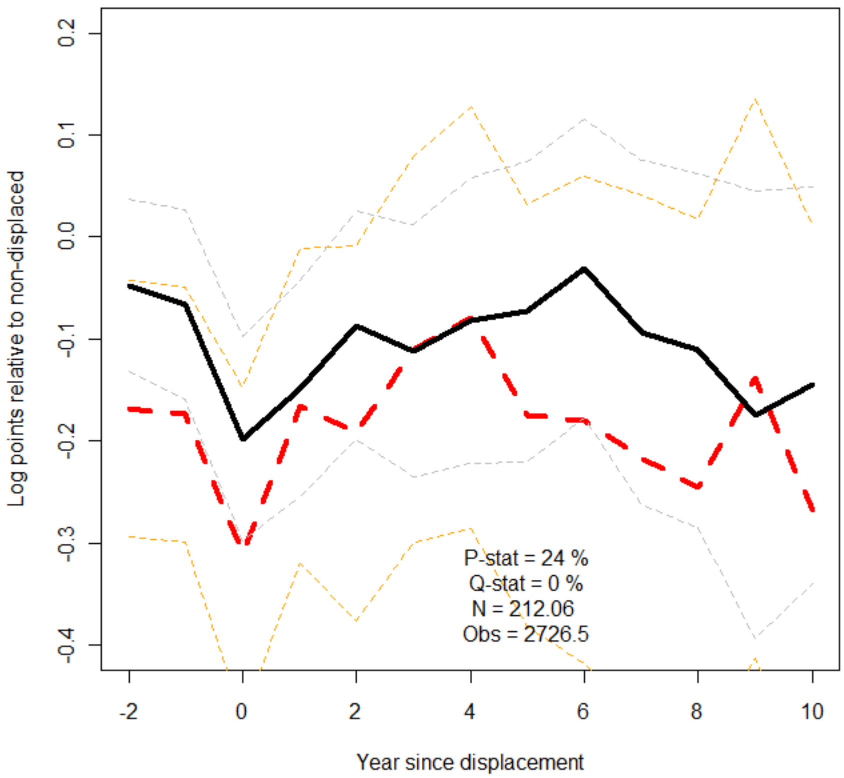


Figure 11: Low task tenure sample; effects of displacement on job task switchers (vs. non-switchers): Coefficients on displacement dummies from estimating a stratified version of equation (10) (stratification on task distance moved) with labor earnings as the dependent variable. This regression used a subset of the analysis sample — which is described in Section 4.2 — which consists of individuals that had low task tenure before displacement. An age quadratic is the only control. The thick red (dotted) line is the marginal effects for non-task switchers and the thick black line is the marginal effects for task switchers who moved the mean distance in task space. Lines representing the confidence intervals ($\pm 2 *$S.E.) of these estimates are also plotted. Standard errors are panel robust.

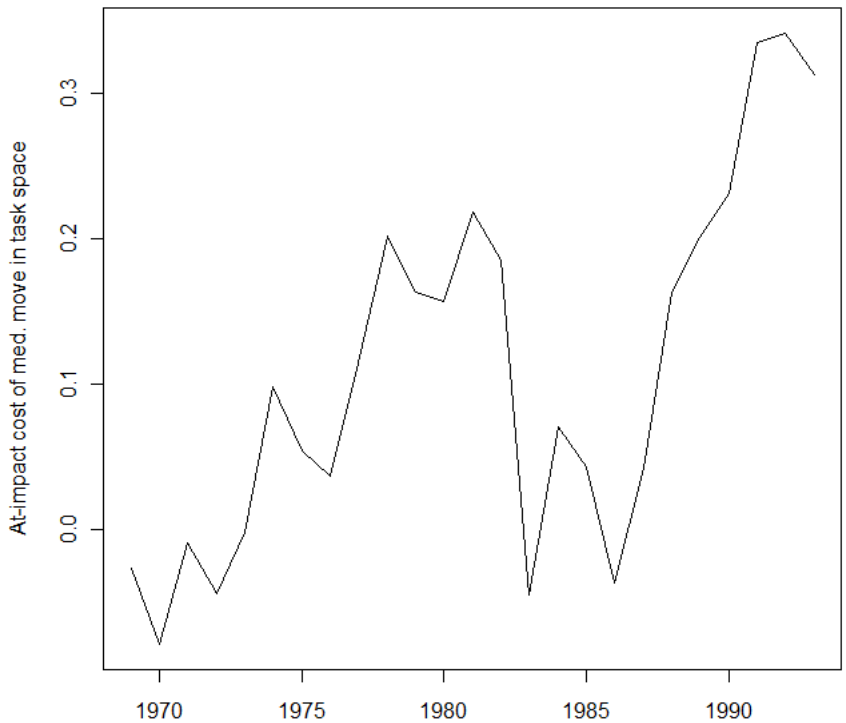


Figure 12: Trend log difference in the at-impact effect of displacement
NOTE: The y-axis is cost of moving a median distance in task space for displaced workers measured in log earnings. Each year is estimated using the "displacement year + 1" coefficient in equation (10) estimated using the sample of displaced workers displaced in that year and the three years before and after.


Figure 13: High-low skill task wage gap and task adjustment costs

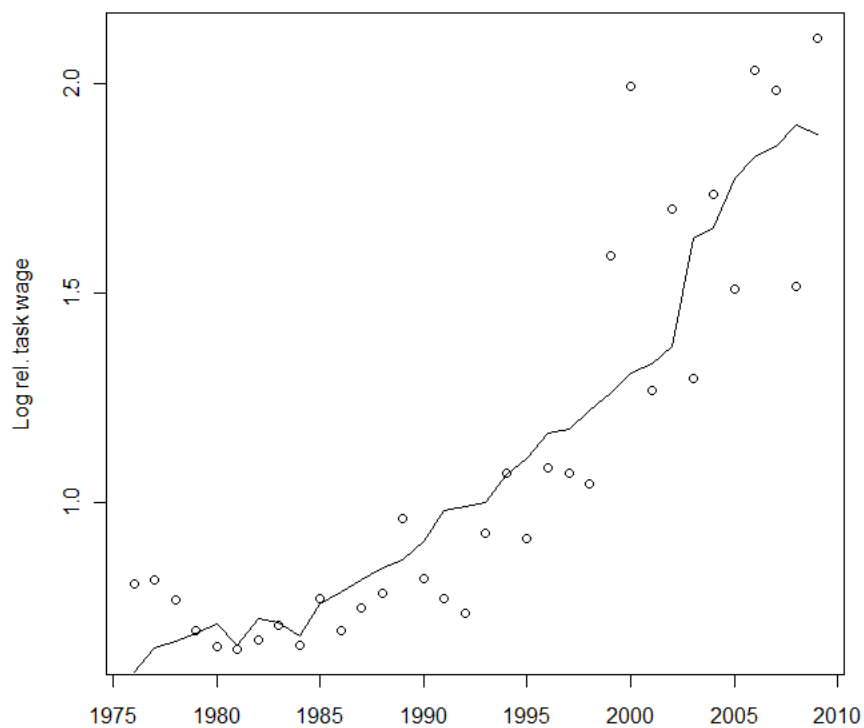


Figure 14: Actual vs. predicted log rel. task wage
NOTE: Dots are actual data and the solid line is the prediction from estimating equation (15).

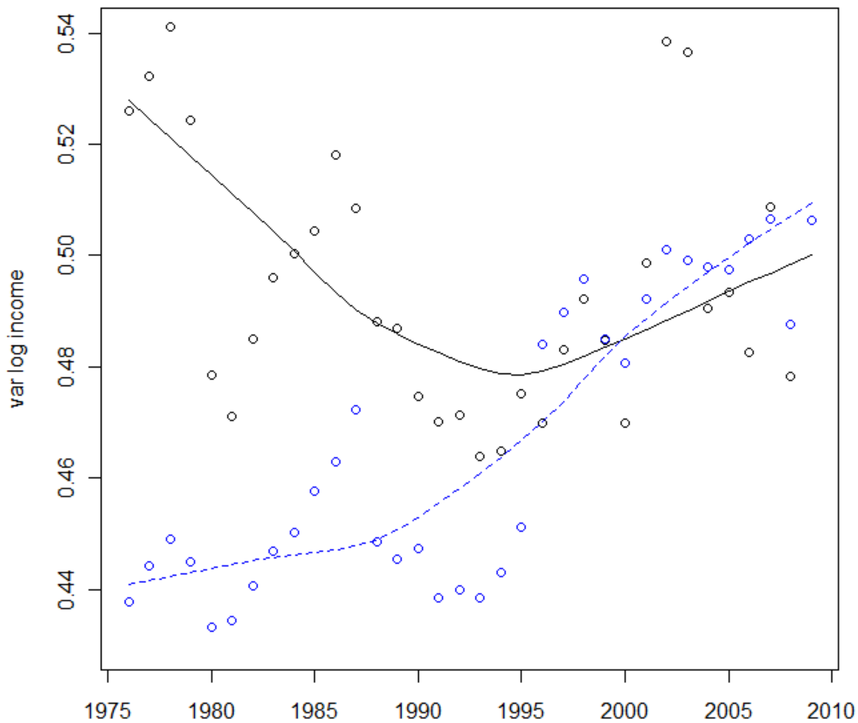


Figure 15: Trends in actual and counter-factual income dispersion
NOTE: A blue dot is the actual variance of log income for each year. The blue dotted line is the lowess fit of those data. A black dot is the counter-factual variance of log income for each year and the black solid line is the lowess fit of those data. To get counter-factual task wages, equation (9) is estimated on 2009 data and employment in all 52 O*Net tasks (see Section 3). For each year, the counter-factual wages are calculated by multiplying the year's task employments by the coefficients from 2009 and then adding the residuals from the task wage regression for the year.