

How Local Are U.S. Labor Markets?: Using an Assignment Model to Forecast the Geographic and Skill Incidence of Local Labor Demand Shocks

Richard K. Mansfield*
University of Colorado-Boulder

PRELIMINARY AND INCOMPLETE
DO NOT CITE WITHOUT AUTHOR'S PERMISSION

October 10, 2018

Abstract

This paper examines how spatial frictions in labor markets differ by worker skill type and firm industry, size, and skill requirements, and how such frictions shape the geographic and skill incidence of alternative local labor demand shocks. LEHD data capturing the near universe of U.S. job transitions from 17 states facilitate the estimation of a rich two-sided assignment model of the labor market featuring thousands of parameters, which is then used to generate simulated forecasts of many alternative local shocks featuring different firm compositions. These forecasts suggest that existing local workers from the targeted census tract for only 0.1% (2.7%) of total welfare (employment) gains from potential 500 job stimulus shocks, with 80% (56%) of welfare (employment) gains accruing to out-of-state workers. This is despite the fact that projected employment rate increases from a typical positive shock are 7 times larger for existing workers in the targeted Census tract than for workers from an adjacent tract, because workers in the target tract are a minuscule share of the national labor market. Further, the projected earnings incidence across local skill groups is quite sensitive to the shock's firm type composition.

*This paper has benefitted greatly from an extraordinary effort by Henry Hyatt, as well as from considerable early assistance from Evan Buntrock, John Abowd, and Erika McEntarfer and valuable comments from Jeronimo Carballo, Terra McKinnish, Brian Cadena, Murat Iyigun, and participants at the NBER Trade and Geography Conference, the IZA World of Labor Conference, and multiple seminars at the University of Colorado-Boulder. This research uses data from the US Census Bureau's Longitudinal Employer-Household Dynamics Program, which was partially supported by NSF grants SES-9978093, SES-0339191, and ITR-0427889; National Institute on Aging grant AG018854; and grants from the Alfred P. Sloan Foundation. All results have been reviewed to ensure that no confidential information is disclosed.

1 Introduction

How are the welfare impacts on workers of local labor demand shocks distributed across geographic space and across skill categories? And how sensitive is this spatial and skill incidence to the firm type composition of the demand shock?

Billions of dollars in local aid are spent each year by state, federal, and local agencies to support city-level or county-level economic development initiatives that seek to enhance labor market opportunities for workers in a particular skill or age class who live or work within the local jurisdiction. These often take the form of local infrastructure spending, discounted loans or subsidies aimed at startup companies, or tax breaks to lure firms to relocate. In order to determine which types of firms or projects to support, federal, state, and local policymakers must predict not only which types of workers from which locations would be directly hired by the tax-supported firms, but also how the resulting ripple effects that operate through vacancy chains and pressure on local wages would indirectly benefit both local and more distant workers.

While a large literature in economics seeks to evaluate the incidence of place-based labor demand policies and shocks, most reduced-form methods focus on quite local impacts, with more distant towns, counties or states either excluded from the sample or used as control groups. Furthermore, by virtue of their focus on particular policies or shocks occurring in one or a small number of locations, these studies are usually ill-equipped to compare the incidence of shocks featuring different demand compositions or to examine differential skill incidence among local and less local areas (due to small samples of workers within a small radius around the shock and/or a lack of detailed data on distant locations).

The primary difficulty is that either evaluating or predicting worker-level welfare incidence across a variety of alternative local labor market shocks requires a spatial equilibrium model that accommodates ripple effects by incorporating the network of spatial linkages among overlapping local labor markets while simultaneously featuring heterogeneity in worker and firm preferences, search costs, and match productivities along a variety of observable dimensions.

Motivated by this challenge, this paper makes two central contributions. First, I develop a theoretically-motivated empirical framework for assessing and forecasting welfare incidence across location-by-demographic group categories of labor demand shocks featuring alternative geographic and firm type compositions (or worker type compositions for labor supply shocks). I do this by adapting to the local labor market setting the two-sided assignment game analyzed originally by Koopmans and Beckmann (1957) and Shapley and Shubik (1972) and whose empirical implications were highlighted in the marriage market context by Choo and Siow (2006). Second, after estimating the parameters of the model, I analyze a large set of model simulations that illustrate several general properties of local labor markets in the United States that effectively create a useful national prior about which

types of workers are most sensitive to which types of local labor demand shocks.

Several key features of Choo and Siow (2006)’s version of the assignment game facilitate these goals. First, it can accommodate multidimensional heterogeneity based on unordered categorical characteristics for agents on both sides of the matching market. In particular, this allows the model to accommodate arbitrary spatial links between different geographic units, including geographic units of both very small and large sizes. It also permits analysis of incidence across demographic groups such as races, age groups, or industries (or combinations thereof) without requiring any hierarchical ordering.

Second, the assignment game requires market clearing, optimizing behavior by all market agents, and explicit payoffs to each agent from each possible job match, making it well-suited for forecasting welfare effects from exogenous shocks. Third, the key parameters of the model (mean relative joint surpluses among matched pairs of workers and firms belonging to observable types) can be identified from a single cross-sectional labor-market transition between origin and destination states, and are sufficient to perform counterfactuals that yield the allocation and impact on payoffs for all players (workers and firms) from any arbitrary change in the composition of labor supply, labor demand, or both.

Finally, these counterfactuals do not require the specification of a more fundamental structural model of utility, firm production, and moving costs, ensuring that none of the heterogeneity present in the transition patterns is lost in paring down to a small number of interpretable structural parameters. The downside to such a “sufficient statistic” approach is that the set of counterfactuals that can be performed is limited to those involving exogenous changes in either the type composition of labor supply and/or demand or composite “joint surplus” parameters. Furthermore, while the heterogeneity on both sides of the labor market can be modeled much more richly than in other structural models, the housing and product markets are not explicitly modeled (though their impact may nonetheless be captured by the estimated surplus parameter through the way they affect job-to-job flows). Thus, the estimated welfare changes only capture “labor-related” welfare changes, and should be thought of as complementary inputs to local policy decisions along with estimates of house price and product price changes.¹

I estimate the model and perform a variety of counterfactual simulations using matched employer-employee data from the Longitudinal Employer-Household Dynamics database on a subset of 17 U.S. states that provide data as of 1993. The data display three key properties that make it suitable for our forecasts. Namely, 1) they capture the (near) universe of job matches from the participating states, mitigating selection problems, 2) they include hundreds of millions of job matches, allowing precise estimates of the large number of parameters necessary to capture the complex two-sided multidimensional heterogeneity, and 3) workers’ establishments are geocoded to the census tract level. These properties,

¹Section 3.4 discusses scenarios in which ignoring the housing market would introduce minimal error into incidence forecasts for the kinds of local labor demand shocks I consider.

when combined, make it feasible to study incidence across worker types at the hyper-local level necessary to make the estimates useful to local policymakers, while still allowing for complex spatial ties between the local area and the surrounding towns, counties, and states that make the estimates useful to state and federal policymakers.

The counterfactual simulations involve firm relocations or stimulus projects that create new job positions in particular U.S. locations (census tracts) featuring alternative combinations of firm size, average pay, and industry supersector. I also consider “natural disaster” simulations akin to a tornado or flood that eliminate a share of all jobs in a particular geographic location.

I find that welfare gains or losses from very local shocks are widely distributed. For example, the simulations suggest that as little as 0.1% of the job-related utility gains and 2.7% of the net employment gains from a stimulus project that creates 500 new jobs in a given census tract accrue to workers already working (or seeking a job) in that tract at the beginning of the year, while 80% (50%) of the job-related utility (employment) gains accrue to workers beginning the year outside the state.

Such geographic dispersion of welfare gains occurs despite the fact that the same simulations suggest that a randomly chosen worker in the targeted census tract is about 7, 25 and 30,000 times more likely to fill one of the new vacancies than a randomly chosen worker in an adjacent tract, 3-5 tracts away, and in a non-adjacent state, respectively. These seemingly inconsistent findings are the result of two mechanisms.

First, since many of the workers likely to join the incoming firms were already employed, so that their transitions generate further openings for others, the share of the stimulus jobs taken by more vs. less local workers dramatically overstates the local concentration of the overall employment and welfare incidence of the labor demand shock. Second (and more importantly), because a single census tract generally only features a few thousand workers, its workforce makes up a very small share of the national labor market. Consequently, even quite substantial and disproportionate welfare gains for the most local workers cannot account for more than a tiny share of the aggregate welfare gains. Indeed, the predicted utility and employment gains for an initially local worker are 2.5 (7), 3.1 (23), and 5.6 (302) times as large as for a randomly chosen worker in an adjacent tract, 3-5 tracts away, and in a non-adjacent state, respectively.

Averaging across simulations, the results suggest that among the most local workers, the utility gains are largest among the initially non-employed (\$896 in money-metric utility units), but are otherwise smaller for initially low-paid (\$396) than initially high-paid workers (\$545), where I use initial earnings as a proxy for worker skill. However, these averages mask substantial heterogeneity in projected impacts across shocks featuring different firm compositions and across sites of simulated shocks. The average money metric utility gains for the same three skill groups are \$1051, \$410, and \$402 for stimuli consisting

of jobs at large, previously low paying firms vs. \$671, \$282, and \$664 for stimuli featuring small, previously high paying firms. I also find that demand shocks consisting of additional jobs at small, low-paying construction firms generate the most locally concentrated employment impact, while small, high-paying information sector jobs generate the smallest local employment impact.

Interestingly, regardless of the firm composition, as the simulated shocks ripple outward, they become less and less skill-biased: predicted differences in welfare (or employment) gains among skill groups converge as we consider workers at initial locations further from the site of the shock.

This paper builds primarily on three literatures. The first consists of evaluations of particular place-based policies or local economic shocks. Most papers in this branch use average wages or employment rates in the targeted location as the outcome of interest, seek to define a control group of alternative locations, and evaluate the policy or shock's impact using a treatment effect framework. This literature is vast, and is thoroughly discussed by survey articles such as Glaeser et al. (2008), Moretti (2010), Kline and Moretti (2013), and Neumark and Simpson (2014).² A few prominent papers in particular stand out.

Autor et al. (2014)'s evaluation of the worker-level impact of China's accession to the WTO is notable for its attention to heterogeneity in incidence across demographic and skill groups. In line with our results, they find that the negative import competition shock particularly affected the cumulative earnings of those with low initial earnings or limited labor force attachment. However, because they consider local variation in the incidence of a national-level shock and use commuting zone industry structure to define shock exposure, their estimates do not provide much guidance on the geographic incidence of a small but geographically concentrated demand shock. Greenstone et al. (2010) is notable for its focus on a relatively large sample of sizable plant openings and relocations (with a known set of counties with losing bids as controls). They show that such plant openings have a considerable impact on countywide employment.

Finally, Busso et al. (2013)'s evaluation of the U.S. empowerment zone system also stands out as one of the few quasi-experimental papers to explicitly evaluate social welfare impact, which they accomplish by deriving a set of sufficient elasticity parameters that can be cleanly identified. Interestingly, they find that while empowerment zones significantly increase wages and employment of zone residents, they do not meaningfully affect rent prices. This suggests that for very local shocks where commuting adjustments play a key role in facilitating the shock response, the impact on rent need not be first-order.³

The paper also contributes to a fast-growing literature on structural spatial equilibrium

²More recent contributions include Gregory (2013), Freedman (2013), and LeGower and Walsh (2017).

³The authors point out that the limited rent price impact might be due to the particularly depressed nature of the targeted locations, which could make them undesirable residential locations (or subject to rent control).

models designed to forecast the incidence of economic shocks across spatially-linked geographic areas. Several recent and contemporaneous papers in particular deserve discussion.

Caliendo et al. (2015) (hereafter CDP) consider the geographic and sectoral incidence of China WTO entry. They develop a full dynamic general equilibrium framework that incorporates input-output linkages in goods markets as well as labor market linkages among a system of 50 U.S. states and 37 countries. They show how counterfactual dynamic equilibrium paths can be evaluated for alternative structural shocks (changes in trade costs, mobility costs, productivities) without estimating all the primitives of the model. Our paper relies on a very similar “sufficient statistics” approach, in that it evaluates the distribution of welfare impacts from demand shocks of alternative compositions without identifying the fundamental utility, production function, and moving cost parameters of the structural model. Like CDP, it relies heavily on the matrix of worker flows for identification of the key parameters that govern these counterfactuals. This paper imposes much less structure on the form of production and utility than CDP, but is also more limited in the set of counterfactuals that can be evaluated. While our paper lacks the explicit housing and product markets modeled in CDP, it features a much richer labor market that highlights existing firm and worker multidimensional heterogeneity and the process by which heterogeneous workers and firms are matched. Our model is thus better able to evaluate differential incidence across skill/demographic groups from labor demand shocks of alternative compositions at a very local level.⁴

Monte et al. (2015) highlights the role of commuting vs. residential mobility in clearing U.S. labor markets across geographic space and in determining the incidence of local labor demand shocks. Like CDP, they also use a trade-theoretic approach to model the joint choice of residential and work location, and incorporate commuting costs, local amenities, and geographic trade costs. Using U.S. counties to define the system of locations, they show that a richer structural model that incorporates the network of labor flows can better predict the heterogeneity in incidence of the million dollar plant openings evaluated by Greenstone et al. (2010). Their model is well-designed to gauge the incidence of shocks to local productivity, housing demand, and commuting technology shocks across locations and across landlords versus residents/workers. However, workers only vary by initial location and unobserved tastes for amenities and locations, and location is the only firm-level attribute that affects labor demand, making it less useful for examining differential incidence across worker categories of different demand shock compositions.⁵

Schmutz and Sidibe (2016) adapt a search and matching framework in the style of McCall (1970) to model the importance of spatial linkages in determining the incidence of local shocks. Based on worker flows between a system of French metropolitan areas, their

⁴The CDP framework could accommodate such features, but at the cost of considerable computational burden.

⁵Due to a lack of micro-level residential data, this paper does not consider whether the new jobs formed by job-to-job transitions involve residential mobility.

results suggest that search frictions play a greater role than moving costs in limiting worker mobility, suggesting the potential promise of efforts to disseminate information about distant jobs. Because firm-level heterogeneity is absent in the job-posting framework they use and workers are ex ante identical beyond initial location, their framework is not designed to evaluate differential shock incidence across demographic groups by shock type.

Each of these papers aggregates locations to at least the county level, leaving considerable room for an analysis of the geographic incidence of very local shocks of the type considered by policymakers in particular towns or cities. Manning and Petrongolo (2017), by contrast, represents the most notable attempt to determine the equilibrium incidence across nearby areas of small scale shocks. Like Schmutz and Sidibe (2016), they propose a search and matching model and derive from it an expression for the equilibrium net outflow of vacancies by area. They fit the predicted geographic distribution of vacancy outflows to data on changes in vacancy stocks from local job search centers in Britain. Like this paper, they simulate the impact on the geographic distribution of unemployment of an exogenous increase in vacancies (new jobs) within particular census wards (similar in size to the census tracts used here). They also find evidence that labor markets are quite local, in the sense that moderate distance to vacancies substantially decreases the probability of an application. Nonetheless, they also find that ripple effects from overlapping markets cause the unemployment incidence to spread widely, with very little of the employment gain accruing to the ward receiving the shock (less than we report here). However, their data sources are the near complement to ours; they observe stocks of reported vacancies and unemployed workers by ward, but have no information on the geographic patterns of either job-to-job or nonemployment-to-job flows.⁶ They also do not consider heterogeneity in incidence across skill types nor by the firm composition of vacancies created, which is a primary goal of this paper.

Finally, this paper also builds upon and draws heavily from the theoretical literature on the identification and estimation of two-sided assignment games. To our knowledge this is the first large-scale labor market application of a two-sided assignment model. The theoretical properties of such assignment games have been well-established for at least a generation⁷. However, the empirical content of the model for contexts in which the universe (or a large random sample) of all market entrants on both sides and their matches can be observed has only recently attracted interest, with Choo and Siow (2006)'s pioneering paper leading to contributions by Chiappori and Salani (2016), Menzel (2015), and Galichon and Salanié (2015), among others. I make three contributions to this theoretical literature.

First, I consider implementation in a context with a very large number of match observations and an even larger number of types on both the supply and demand side. I address

⁶See Nimczik (2017) for an effort to define the boundaries of local labor markets based on job-to-job flows rather than commuting patterns.

⁷See Koopmans and Beckmann (1957), Shapley and Shubik (1972), Roth and Sotomayor (1992), and Sattinger (1993) among others

this problem by introducing a smoothing procedure designed to aggregate matching patterns across “nearby” match types without smoothing away the heterogeneity the model is designed to highlight.

Second, I consider the limits to identification in a context where the number of unmatched partners of each type is either unobserved or only observed on one side of the market: while nonemployment may be inferred with reasonable accuracy in the LEHD data, unfilled vacancies are absent. The existing identification strategies employed by Choo and Siow (2006) and Menzel (2015), among others, rely heavily on observing the number of singles on both sides. In particular, I discuss conditions under which ignoring unmatched partners would not affect the incidence of policy interventions among originally matched agents.

Third, I investigate the impact of relaxing the assumption that agents on each side of the market only have preferences for agent types on the opposite side, rather than particular agents (so that unobserved heterogeneity at the job match level is introduced). I show that such essential match-level heterogeneity does not undermine the ability to forecast welfare incidence from shocks to the composition and level of labor supply or demand, though it does complicate the use of observed transfers to separately identify the worker and firm pre-transfer surplus values from alternative matches.

The rest of the paper proceeds as follows. Section 2 describes the two-sided assignment game that forms the theoretical basis for the empirical analysis. Section 3 illustrates how to apply the insights of Choo and Siow (2006) to the context of labor market transitions to identify a set of joint surplus parameters that are sufficient to perform counterfactual simulations of labor demand shock incidence. Section IV describes the LEHD database. Section V describes sample selection and the smoothing procedure used to eliminate disclosure risk and minimize the sparsity of the large-scale transition matrix whose entries determine the surplus parameters. In addition, Section V also provides detail about the particular specifications of labor demand shocks that I simulate, the procedure used to perform the simulations, and the methods used to aggregate the resulting counterfactual allocations of workers to positions into interpretable statistics that effectively characterize variation in shock incidence. Section VI presents the main findings, and Section VII concludes the paper.

2 Model

2.1 Model Overview

In this section I model the evolution of the labor market across time periods as a sequence of static cooperative matching games played by workers and firms. The model is based on Choo and Siow (2006)’s model of the marriage matching market, but introduces a number

of features and extensions necessary to adapt their model to a labor market setting. The exposition of the model closely mirrors Galichon and Salanié (2015), which generalizes Choo and Siow (2006). Section 2.2 lays out the basics of the matching game. Section 2.3 describes how the workers and positions (the game’s agents) and the job matches that determine the game’s payoffs are aggregated to types and groups, respectively. Section 3 imposes additional structure on the model that facilitates the identification and estimation of the underlying group-level match surpluses that determine the frequencies of particular kinds of job transitions. Section 3.2 shows how these estimated match surpluses can be used to construct counterfactual simulations capturing the incidence of local labor supply and demand shocks of varying worker and firm compositions.

2.2 Defining the Assignment Game

Suppose that in a given year y there are I potential workers in the labor market, with the set of individual workers denoted \mathcal{I} . Each worker begins the year in a job match with a position $j(i)$ at establishment $m(j(i))$, determined in year $y-1$, from the set of possible firm positions \mathcal{J} . Let positions featuring $m(j) = 0$ represent nonemployment so that positing an initial “job” match for each worker is without loss of generality.

Suppose that the value to worker i currently at position j of accepting a position k the following year is given by:

$$U(i, j, k) = f^i(A(k)) + w_{ik} - 1(m(j) \neq m(k))c^i(m(j), m(k)) + \beta E_{k'}[U(i, k, k')] \quad (1)$$

In equation (1), each worker’s flow utility in the year following the job transition is assumed to be captured by a quasi-linear money-metric utility function $u_i(w, A(k)) = f^i(A(k)) + w_{ik}$. $A(k)$ captures a vector of non-pecuniary amenities offered by position k , and $f^i(*)$ reflects worker i ’s valuation of these amenities. These might include injury risk, schedule flexibility, or the desirability of the geographic location of the position, and can be differently valued by different individuals. w_{ik} captures worker i ’s yearly earnings at position k (determined in equilibrium).⁸ $c^i(m(j), m(k))$ captures the cost to the worker of relocating from a position at establishment $m(j)$ to one at establishment $m(k)$, which could include search costs, moving costs (which may be related to the geographic distance between establishments $m(j)$ and $m(k)$), or training costs borne by the worker.⁹ In this specification, these costs are only

⁸Since I have data on annual earnings but not wages or hours, for simplicity I assume that the hours associated with a job match are fixed by contract and common across positions for a given worker, and focus exclusively on earnings.

⁹The traditional assignment game does not feature any stochastic search frictions, so that each agent might in principle match with any agent on the opposite side of the market. However, Menzel (2015) shows that one can introduce a probability $r(i, k)$ that i and k meet that is independent of other payoff determinants, assign the utility from the match to $-\infty$ if the pair does not meet, and use these alternative payoffs to determine the stable matching. Alternatively, search costs might be modeled as a deterministic cost that must be paid by an agent to an intermediary (e.g. a headhunter or a job matching website) to

borne by the worker if the worker changes establishments. $\beta E_{k'}[U(i, k, k')]$ captures the discounted value of the worker's job search in the following year, given that he/she will begin the year at position k , which affects the future value by changing the distribution of job switching/relocation costs for alternative positions. By combining the non-earnings components of the worker's valuation into a single index π_{ijk}^i , one can rewrite the value function $U(i, j, k)$ as:

$$U(i, j, k) = \pi_{ijk}^i + w_{ik} \quad (2)$$

From this point forward, because j is fully determined by i in the static assignment game I consider, I will frequently suppress j and replace $U(i, j, k)$ with $U(i, k)$, and will do the same with other functions that depend on j , except where necessary to illustrate the role of workers' origin job assignments in determining the reallocation of workers across job matches.

On the other side of the market there are K potential positions at establishments that seek workers in year y that make up the set \mathcal{K} . Note that the intersection of the sets \mathcal{K} and \mathcal{J} may be quite large, so that many of the end-of-period positions in \mathcal{K} can potentially be "filled" by simply continuing a job match that already exists. We assume that each establishment makes independent hiring decisions for each position, so that we can model the preferences of positions over individual workers rather than modeling establishment preferences over collections of workers.¹⁰ Let the value of hiring (or retaining) a given worker i to a particular position k in establishment $m(k)$ be given by:

$$V(i, j, k) = R_k(i) - w_{ik} - 1(m(j(i)) \neq m(k))c^k(i, m(k)) + \beta E_{i'}[V(i', j(i'), k)] \quad (3)$$

Here, $R_k(i)$ captures the contribution of worker i to firm $m(k)$'s revenue in the coming year and w_{ik} reflects the annual earnings paid to worker i . $c^k(i, m(k))$ captures any training costs, search costs, and recruiting costs borne by firm $m(k)$ in hiring worker i , which are only incurred if worker i is not already filling a position at establishment $m(k)$, so that $m(j) \neq m(k)$.¹¹ $\beta E_{i'}[V(i', j(i'), k)]$ represents the discounted future value of position k . While not made explicit by the notation, this future value incorporates the fact that worker

reveal/allow contact with particular agents on the opposite side of the market (without which they cannot be found).

¹⁰One justification for treating positions as independent is that there are nontrivial costs of coordinating multiple independent hires/retention decisions that outweigh the gains from better exploiting the complementarities that exist in the production process. See Roth and Sotomayor (1992) for a detailed analysis of how the properties of the assignment game change when establishments have preferences over collections of workers.

¹¹Note that imposing that the training/recruiting cost occurs at the establishment rather than the position level suggests some dependence in hiring decisions across positions within an establishment, since a worker hired for one position would not require a second search/recruiting cost if he/she transitioned to another position within the establishment. This slight abuse of notation could be rectified by adding an extra term in the continuation value internalizing this externality, but we omit this term to keep the specification simple. While there may also exist position-level switching costs, our data do not allow us to identify transitions among different positions within the same establishment.

i would begin the following year at position k , so that retaining i would not require further recruiting/training costs. As with the worker's value, we can also suppress dependence on j , form an index of the non-pecuniary components of the firm's valuation, π_{ik}^k , and rewrite the value function as:

$$V(i, k) = \pi_{ik}^k - w_{ik} \quad (4)$$

Using the simplified worker and position value functions, we can define the joint surplus from the transition of worker i to position k as the sum of the worker and position valuations of the transition:

$$\pi_{ik} \equiv U(i, k) + V(i, k) = \pi_{ik}^i + \pi_{ik}^k \quad (5)$$

Note that by assuming money-metric value functions that impose additive separability of current worker earnings in both the worker's and firm's payoffs, we have ensured that the assignment model exhibits transferable utility, with earnings representing the transfer. Thus, once the payoffs are written in this form, the game has the exact structure of the classic assignment game analyzed by Koopmans and Beckmann (1957) and Shapley and Shubik (1972).

Importantly, note that while particular subcomponents of the worker and firm value functions were specified in (1) and (3), this was done purely to provide intuition for the deeper structural parameters that might enter into the surplus components π_{ik}^i and π_{ik}^k , and thus to better gauge when the assumptions necessary for valid counterfactuals laid out in section 3.2 below are likely to be satisfied. Any value function specifications in which current worker earnings are additively separable from other joint payoff determinants will suffice; I show below that the researcher need not specify any of the fundamental parameters or the functions governing their links to payoffs in order to construct the counterfactual simulations that form the primary contribution of the paper.

A matching or market-wide transition in this labor market is an $I \times K$ transition matrix μ such that $\mu_{i,k} = 1$ if worker i matches with position k at the end of the period, and 0 otherwise. As in Galichon and Salanié (2015), I focus on stable matchings, which have the property that a division of joint surplus exists in each job match such that no currently unmatched worker-position pair can generate a joint surplus that features a division among the pair that makes both the worker and position strictly better off than they are under the proposed matching. Shapley and Shubik (1972) showed that the set of stable matchings coincides with both the core of the assignment game and with the set of competitive equilibria from a decentralized labor market. Furthermore, they show that in the presence of transferable utility there will exist a unique assignment (or, equivalently, competitive equilibrium allocation) of origin job matches to destination job matches as long as preferences are strict on both sides of the market. This equilibrium allocation/stable assignment will maximize the aggregate surplus, and thus can be found by solving a linear programming

problem.¹²

Equivalently, the unique stable assignment can also be found by solving the dual problem: identifying a set of worker discounted utility values $\{r_i\}$ and firm discounted profit values $\{q_k\}$ that minimize the total “cost” of all workers and firms, $\sum_{i \in \mathcal{I}} r_i + \sum_{k \in \mathcal{K}} q_k$, subject to the constraint that these values cannot violate the underlying joint surplus values: $r_i + q_k \geq \pi_{ik} \forall (i, k)$. This dual problem yields the following conditions that define the optimal assignment (Galichon and Salanié (2015)):

$$\mu_{ik} = 1 \text{ iff } k \in \arg \max_{k \in \mathcal{K} \cup 0} \pi_{ik} - q_k \text{ and } i \in \arg \max_{i \in \mathcal{I} \cup 0} \pi_{ik} - r_i \quad (6)$$

I aggregate these conditions in the next section to deliver identification of aggregate “type”-level surpluses. Moreover, Koopmans and Beckmann (1957) point out that only the set $\{r_i\}$ (or alternatively $\{q_k\}$) is necessary to construct the stable assignment, since the maximizing choices for the position side of the market in (6) determine the full set of job matches, and the maximized value for each position k is its q_k value.

Finally, given optimal (or equilibrium) worker and firm payoffs $\{r_i\}$ and $\{q_k\}$ from the dual solution, Shapley and Shubik (1972) show how to decentralize this optimal assignment via a set of earnings transfers w_{ik} that appear in the worker and firm value functions above:

$$w_{ik} = \pi_{ik}^k - q_k \quad (7)$$

Because $r_i + q_k = \pi_{ik} \equiv \pi_{ik}^i + \pi_{ik}^k$ for any pair (i, k) that is matched in the unique stable match, this also implies that:

$$w_{ik} = r_i - \pi_{ik}^i \quad (8)$$

Then the conditions (6) can be rewritten as the standard requirements that worker and firm choices must be utility- and profit-maximizing, respectively:

$$\mu_{ik} = 1 \text{ iff } k \in \arg \max_{k \in \mathcal{K} \cup 0} \pi_{ik}^i + w_{ik} \text{ and } i \in \arg \max_{i \in \mathcal{I} \cup 0} \pi_{ik}^k - w_{ik} \quad (9)$$

This shows that the market-clearing earnings amounts will in general be specific to worker-position pairs (i, k) . By contrast, the market-clearing utilities r_i and profit contributions q_k will be worker-specific and position-specific, respectively, which is a property that will be exploited below. Importantly, while the stable assignment μ_{ik} is generally unique, the reservation values and wages are not: all $\{r_i\}$ values can generally be shifted either up or down by a small increment δ (with offsetting decreases or increases for $\{q_k\}$) without violating any of the stability conditions.¹³ The exact reservation values/wages that emerge

¹²The joint surplus is given by $\sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}} \mu_{i,k} \pi_{ik} + \sum_{i \in \mathcal{I}; \mu_{i,k}=0 \forall k} \mu_{i,0} \pi_{i0} + \sum_{k \in \mathcal{K}; \mu_{i,k}=0 \forall i} \mu_{0,k} \pi_{0k}$. Here, π_{i0} and π_{0k} denote worker i 's payoff from nonemployment and firm k 's payoff from an unfilled vacancy. This sum is maximized subject to the constraints that each worker and firm can be matched to at most one counterpart: $\sum_{i'} \mu_{i',k} \leq 1 \forall k \in \mathcal{K}$ and $\sum_{k'} \mu_{i,k'} \leq 1 \forall i \in \mathcal{I}$

¹³Thus, when generating counterfactual allocations in the empirical work below, we will often normalize

depend on the particular process by which the decentralized labor market converges.

While the analysis does not require taking a stand on a particular earnings determination process, it is nonetheless illuminating to present one candidate process: a simultaneous ascending auction. In such an auction, positions are bidding on all workers simultaneously. Each position may only be the highest bidder for a single worker (or for none, if it chooses to remain vacant). Workers may set reservation utilities which will vary based on the value different workers place on remaining nonemployed for a year. The position k that bids the highest discounted utility r_i to a worker i retains the worker and pays the worker an annual earnings amount w_{ik} that, when combined with the non-pecuniary component π_{ik}^i , is sufficient to ensure the worker's promised valuation $U_{ik} = r_i$. The auction ends when no position wishes to outbid any other position for a worker. Some workers may remain nonemployed, and some positions may remain unfilled.

Importantly, under the quasi-linear utility specification above, even though positions are bidding values of a one period commitment U_{ik} (which include continuation values), they each start at different baseline levels of π_{ik}^i , and changes in bidding can always take the form of earnings increases. Thus, when solving for changes in the stable assignment of workers to firms following shocks to labor demand composition below, we can compute the changes in reservation utilities r_i that clear the market, and we may scale these changes in terms of annual earnings gains (though in some cases workers achieving utility gains will take an earnings cut to work at a firm offering non-pecuniary values that more than offset the earnings decrease).

2.3 Modeling the Match Surpluses

Part of the joint transition surplus π_{ik} for worker-position pair (i, k) from transition $(i, j(i))$ to (i, k) (henceforth denoted (i, k)) is likely to be common to any transition $(i', j(i'))$ to (i', k') that shares certain salient characteristics of the worker, positions, origin or destination job matches, or even transition. For example, positions at larger firms may face smaller per-position costs of recruiting distant workers due to economies of scale; highly skilled workers may generate larger surplus at positions whose output is particularly sensitive to worker skill. Thus, we assign each potential transition (i, k) to one of a set of mutually exclusive groups $g \in \mathcal{G}$ (with $G \equiv |\mathcal{G}|$), and use the notation $g(i, k)$ to denote the group to which transition (i, k) has been assigned. Importantly, these groups are always defined by a combination of *observable* characteristics of the worker i , the origin job $j(i)$ or job match $(i, j(i))$, the destination position k , or even the transition (i, k) . The characteristics that define the set of transition groups should be chosen to capture as comprehensively but parsimoniously as possible the underlying (structural) preferences, productivities, mov-

changes in utility values for one type of worker to be zero, and analyze relative changes that are better identified.

ing costs, and geographic search costs that enter into (1) and (3) and thus determine the relative desirability of the job match for both the worker and the position.

Some subset of these observed characteristics may only relate to the worker i , the origin position $j(i)$, or the worker-origin position job match $(i, j(i))$, and will be common to all destination positions $k' \in \mathcal{K}$. I use this subset of characteristics to assign each worker-origin position job match $(i, j(i))$ to an origin type $o \in \mathcal{O}$, and use the notation $o(i)$ or $o(i, j(i))$ to denote the origin type to which $(i, j(i))$ has been assigned. In the empirical work below, the origin types are defined by unique combinations of the following pair of characteristics: the geographic location (either census tract, public-use micro area PUMA, or U.S. state) of the establishment j at which worker i works in the origin period and the skill level of worker i (proxied by the national earnings quintile associated with i 's earnings while at position j (in year $y - 1$)).¹⁴

Analogously, another (mutually exclusive) subset of the observed characteristics defining the transition group g may only characterize the destination position (i.e. they are common to all origin workers i). This subset is used to assign each destination position $k \in \mathcal{K}$ to a destination type $d \in \mathcal{D}$. In our empirical work, the destination types are defined by unique combinations of the following characteristics: the geographic location (U.S. census tract, PUMA, or state) of position k 's establishment, the establishment size quartile of the establishment associated with position k (in the origin year) in the national establishment-level employment distribution, the quartile of average worker earnings at establishment $m(k)$ (in the origin year) in the national establishment-level average earnings distribution (intended to proxy for the average required skill level of establishment $m(k)$), and the industry supersector of the firm associated with position k .¹⁵

Finally, let $z(i, j(i), k) \equiv z(i, k)$ capture the remaining subset of characteristics defining the transition group that depend on both $(i, j(i))$ and k . In the empirical work below, the single z characteristic will be an indicator for whether the “transition” represents continued employment at the same establishment, $1(m(k) = m(j(i)))$, and is intended to capture the fact that search, recruiting, and training costs do not have to be repaid by existing workers (as reflected in the terms $1(m(j) \neq m(k))c^i(m(j), m(k))$ and $1(m(j) \neq m(k))c^k(i, m(k))$ in the value functions (1) and (3) above). This allows the placement of job stayers and

¹⁴Since the goal is to characterize the geographic scope of workers' searches for jobs and firms' searches for employees, we would ideally use residential location to define the origin job type and establishment location to define the destination job type. However, we do not have access to the residential location of the worker, so we use establishment location in year $y - 1$ as a proxy for residential location in year $y - 1$.

¹⁵Note that we interpret the earnings category of i in year $y - 1$ as a proxy for worker i 's skill, but are interpreting the average earnings category of establishment $m(k)$ in year y as a proxy for the skill requirements of position k . This seeming inconsistency can be rationalized if we assume that the worker, once hired, develops the skills required by the job by the end of the year (perhaps incurring some training costs that could be borne by either the worker or the firm and thus affect the joint surplus π_{ik}). The assignment model also suggests that past equilibrium transfers will generally reflect both the skill requirements of the firm and the skill requirements of the worker, so that they may be effective proxies for both in the absence of other information.

job movers into different groups, which in turn allows establishments to retain existing employees at different rates than they hire other local workers (important for predicting which workers ultimately accept newly created jobs).

Thus, without loss of generality we can rewrite the mapping $g(i, k)$ as $g(o(i), d(k), z(i, k)) \equiv g(o, d, z)$. Importantly, while knowledge of g is sufficient to recover both o and d , knowledge of o and d need not uniquely identify the group g (due to the presence of z). In a slight abuse of notation, I will sometimes use $o(g) = o(g(i, k)) = o(i)$ to refer to the origin type associated with group g , and I will use $d(g) = d(g(i, k)) = d(k)$ to refer to the destination type associated with group g .

Given these definitions, we can decompose the transition surplus π_{ik} into the part that is common to all transitions classified as group $g(i, k)$, denoted θ_g , and an idiosyncratic component ϵ_{ik} specific to the particular transition (i, k) :

$$\pi_{ik} = \theta_{g(i,k)} + \sigma\epsilon_{ik} \equiv \theta_{odz} + \sigma\epsilon_{ik} \quad (10)$$

ϵ_{ik} might reflect, for example, the low psychic costs of a particular worker who is moving back to the location where his family lives, or perhaps particular skill requirements of position k that worker i uniquely possesses. We assume below that ϵ_{ik} is independent and identically distributed across all alternative matches (i, k') and $(i', k) \in \mathcal{I} \times \mathcal{K}$ and follows a Type 1 extreme value distribution. σ is a scaling parameter that captures the relative importance of idiosyncratic components of the matching surplus compared to components that are common among all transitions classified into the same group g in determining the variation in match surpluses across potential pairs $\{(i, k) \in \mathcal{I} \times \mathcal{K}\}$.

The goal is to use the observed matching μ to recover the set of group mean surplus values $\{\theta_g\}$. As Galichon and Salanié (2015) emphasize, one way to do this is to impose further structure on the production, utility, search cost, and recruiting cost functions in (1) and (3), so that $\theta_g \rightarrow \theta_g(\lambda_1)$ for some smaller set of structural parameters λ_1 , with the distribution of ϵ_{ik} depending on a second parameter set λ_2 . Maximum likelihood can then be used to relate the observed match μ to the parameters of the model.

Driven by a combination of computational considerations and an interest in being agnostic about the various structural functions that underlie $\{\theta_g\}$, I follow Choo and Siow (2006) and leave the set $\{\theta_g\}$ unrestricted, achieving identification instead by assuming that ϵ_{ik} is distributed Type 1 extreme value. The aggregated assignment model just laid out differs in three important ways from that of Choo and Siow (2006). First, we allow for the possibility that part of the match surplus is truly idiosyncratic: the combined surplus from two transitions (ik) and $(i'k')$ would be altered if the two workers swapped destination positions, even if i and i' are both associated with the same origin type ($o(i) = o(i')$) and k and k' are both associated with the same destination type ($d(k) = d(k')$). Given the coarseness of the origin and destination types in our empirical work, such within-type-combination

heterogeneity in match quality is very likely to exist in the labor market. Allowing such match-level heterogeneity, however, comes at a cost: as discussed in Section 3.2 and Appendix A4, we forfeit a straightforward way to use observed transfers to separate the group mean surplus θ_g into group-level worker and firm subcomponents θ_g^l and θ_g^f analogous to the transition-level surplus components π_{ik}^i and π_{ik}^k defined above. Fortunately, this further decomposition is not necessary to perform an important class of counterfactual simulations (discussed further below).

Second, we allow for multiple groups within origin/destination type combination (o, d) , so that $\theta_{od} \rightarrow \theta_{odz}$. In our setting, allowing for job switching costs even within the same local labor market is essential for making accurate predictions about how much reallocation a given shift in labor demand or supply will cause.¹⁶

Third, we allow for a separate parameter, σ , that captures the relative importance of idiosyncratic match-level factors versus group-level factors in determining the overall surplus from a job transition. As we will see below, the introduction of σ will not change the unique stable job assignments in our counterfactual simulations, but it will play a key role in determining the size of changes in offered utility values r_i for particular workers in particular locations that are necessary to facilitate the reallocation that yields the stable assignment.

3 Identification

3.1 Identification of the Set of Group-Level Match Surpluses $\{\theta_g\}$

Recall from section 2.2 that a necessary condition for a matching μ to be stable (and thus sustainable as a competitive equilibrium) is that there exists a set of worker values $\{r_i\}$ such that $\mu_{ik} = 1$ implies that $i \in \arg \max_{i \in \mathcal{I}} \pi_{ik} - r_i$ for any potential match $(i, k) \in \mathcal{I} \times \mathcal{K}$.

Given the Type 1 extreme value assumption for ϵ_{ik} in equation (10), Decker et al. (2013) show that the probability that an existing position k is filled by hiring (or continuing to employ) i is given by:

$$P(i|k) = \frac{e^{\frac{\theta_g - r_i}{\sigma}}}{\sum_{i' \in \mathcal{I}} e^{\frac{\theta_{g'} - r_{i'}}{\sigma}}} \quad (11)$$

where we have suppressed the dependence of g and g' on (i, k) and (i', k) , respectively. We can then use equation (11) to derive an expression for the probability that a randomly chosen position associated with destination type d hires a worker whose transition to the

¹⁶In their marriage market setting, one could classify end-of-year matches among a given male type/female type combination (o, d) with an additional variable z capturing whether the match is a continuing marriage or a new marriage, or by the allocation of household work within the marriage. Data on attributes of the match (beyond the attributes of the agents creating the match) often exist that can enrich our understanding of how the terms of matches are negotiated or how they succeed/fail.

position is classified into group g :

$$\begin{aligned}
P(g|d) &= \sum_{k \in d} P(g|d, k)P(k|d) = \frac{1}{|d|} \sum_{k \in d} P(g|k) \\
&= \frac{1}{|d|} \sum_{k \in d} \sum_{i: g(i, k)=g} P(i|k) \\
&= \frac{1}{|d|} \sum_{k \in d} \sum_{i: g(i, k)=g} \frac{e^{\frac{\theta_g - r_i}{\sigma}}}{\sum_{i' \in \mathcal{I}} e^{\frac{\theta_{g'} - r_{i'}}{\sigma}}} \\
&= \frac{1}{|d|} \sum_{k \in d} \frac{(e^{\frac{\theta_g}{\sigma}})(\sum_{i: g(i, k)=g} e^{\frac{-r_i}{\sigma}})}{\sum_{i' \in \mathcal{I}} e^{\frac{\theta_{g'} - r_{i'}}{\sigma}}} \tag{12}
\end{aligned}$$

Next, we make two assumptions that allow us to express this conditional probability exclusively in terms of the group g and the types o, d . First, Assumption 1 imposes that the mean exponentiated worker utility values $e^{\frac{-r_i}{\sigma}}$ vary minimally across groups g featuring the same origin type $o(g)$. Given the characteristics used to define o and g in our application, this states that existing employees (potential stayers) and non-employees of each establishment have approximately the same mean value of r_i among workers whose initial jobs were in the same local area and pay category. In other words, the payoffs that workers in the same skill class require in equilibrium will not differ systematically across establishments within a small local area. This becomes a better approximation the more characteristics (such as occupation or education) are used to define an origin type $o(i)$. To formalize this assumption, recall that the only characteristic z that distinguishes transition groups featuring the same combination of origin and destination types (o, d) is an indicator for whether the worker i was already employed by k in the previous period, so that a given (o, d) pair contains at most two groups, potential stayers and potential new hires. We can thus write:

$$\text{Assumption 1: } \frac{1}{|g_k|} \sum_{i: g(i, k)=g} e^{-\frac{r_i}{\sigma}} \approx \frac{1}{|o|} \sum_{i: o(i)=o(g)} e^{-\frac{r_i}{\sigma}} = C_{o(g)} \quad \forall (g, k) \tag{13}$$

where $|o|$ and $|g_k|$ denote, respectively, the number of workers classified as origin type o and the number of workers whose transition would be classified as group g (either stayers or new hires among those in o) if hired by position k (a subset of the workers in $o(g)$). C_o denotes the mean value of $e^{-\frac{r_i}{\sigma}}$ for a given origin group o .

Second, Assumption 2 imposes that the number of potential stayers vs. new hires among workers from a given origin type o is common across establishments assigned to the same destination type d . In our context, this means that establishments in the same geographic area that are in the same industry supersector and same establishment size and establishment average pay categories have roughly the same number and skill composition

of employees. Formally:

$$\text{Assumption 2: } P(z(g)|o(g), k) \approx P(z(g)|o(g), d(k)) \forall k, \forall g \quad (14)$$

Letting $f(o)$ denote the share of all workers assigned to origin type o , this implies:

$$|g_k| \equiv P(z(g)|o(g), k)f(o(g))|\mathcal{I}| \approx P(z(g)|o(g), d(k))f(o(g))|\mathcal{I}| \quad (15)$$

These assumptions are necessary because the aggregate mean of $e^{-\frac{r_i}{\sigma}}$, a non-linear function of a random variable, depends on its entire distribution. Essentially, the probability of filling a position with an existing employee depends on how many employees one already has, so that the group average depends on the firm size distribution among firms who are at risk of creating a transition that could be classified into g . I am essentially hoping that Jensen's inequality is close to equality ($f(E[X]) \approx E[f(X)]$) after conditioning on the characteristics that define our origin and destination types (most notably establishment size category). Note that the set of conditional probabilities $P(z(g)|o, d) \equiv P(z(i, k) = z(g)|o, d)$ can be assigned/estimated for each group g prior to the rest of estimation using the average across all positions in group g of the fraction of candidates for the position that are existing employees.

Combined, these two assumptions imply that

$$\sum_{i:g(i,k)=g} e^{-\frac{r_i}{\sigma}} \approx P(z(g)|o(g), d(k))f(o(g))|\mathcal{I}|C_{o(g)}. \quad (16)$$

Applying these assumptions to the last expression in (12), we obtain:¹⁷

$$P(g|d) = \sum_{k \in d} \left(\frac{1}{|d|} \right) \frac{e^{\frac{\theta_g}{\sigma}} \sum_{i: g(i,k)=g} e^{-\frac{r_i}{\sigma}}}{\sum_{i' \in \mathcal{I}} e^{\frac{\theta_{g'} - r_{i'}}{\sigma}}} \quad (17)$$

$$= \sum_{k \in d} \left(\frac{1}{|d|} \right) \frac{e^{\frac{\theta_g}{\sigma}} \sum_{i: g(i,k)=g} e^{-\frac{r_i}{\sigma}}}{\sum_{o' \in \mathcal{O}} \sum_{g' \in (o,d)} \sum_{i': g(i',k)=g'} e^{\frac{\theta_{g'} - r_{i'}}{\sigma}}} \quad (18)$$

$$= \sum_{k \in d} \left(\frac{1}{|d|} \right) \frac{e^{\frac{\theta_g}{\sigma}} P(z(g)|o, d) f(o) |\mathcal{I}| C_o}{\sum_{o' \in \mathcal{O}} \sum_{g' \in (o,d)} e^{\frac{\theta_{g'}}{\sigma}} P(z(g')|o', d) f(o') |\mathcal{I}| C_{o'}} \quad (19)$$

$$= \frac{e^{\frac{\theta_g}{\sigma}} P(z(g)|o, d) f(o) |\mathcal{I}| C_o}{\sum_{o' \in \mathcal{O}} \sum_{g' \in (o,d)} e^{\frac{\theta_{g'}}{\sigma}} P(z(g')|o', d) f(o') |\mathcal{I}| C_{o'}} \sum_{k \in d} \left(\frac{1}{|d|} \right) \quad (20)$$

$$= \frac{e^{\frac{\theta_g}{\sigma}} P(z(g)|o, d) f(o) C_o}{\sum_{o' \in \mathcal{O}} \sum_{g' \in (o,d)} e^{\frac{\theta_{g'}}{\sigma}} P(z(g')|o', d) f(o') C_{o'}} \quad (21)$$

Let $\hat{\mu}$ denote an observed empirical matching. Since each observed transition can be assigned to a unique group g , we can easily aggregate the empirical matching into a group-level empirical distribution of transitions. Specifically, we let \hat{P}_g denote the fraction of all observed transitions that are assigned to group g ; $\hat{P}_g \equiv \frac{1}{|\mathcal{I}|} \sum_{(i,k) \in \mathcal{I} \times \mathcal{K}} \hat{\mu}_{ik} 1(g(i,k) = g)$. Similarly, $\hat{f}(o)$ denotes the empirical fraction of all job transitions whose worker can be classified as type o : $\hat{f}(o) = \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} 1(o(i) = o)$. Finally, $\hat{h}(d)$ denotes the empirical fraction of all job transitions whose destination position can be classified as type d : $\hat{h}(d) = \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} 1(d(k) = d)$.¹⁸ Based on these definitions, we can estimate the (year-specific) conditional choice probability $P(g|d)$ by simply calculating the observed fraction of destination positions classified as type d that were filled via transitions assigned to group g : $\hat{P}(g|d) = \frac{\hat{P}_g}{\hat{h}(d)}$. Consequently, as the number of observed transitions gets large, each member of the set of empirical conditional choice probabilities $\{\hat{P}(g|d)\}$ should converge to the corresponding expression in (21).

We are now ready to investigate the extent to which these assumptions, combined with the observed empirical choice probabilities $\{\hat{P}(g|d)\}$, can inform us about the mean match surplus values $\{\theta_g\}$. Consider the log odds between two conditional choice probabilities involving an (arbitrarily chosen) destination type d_1 and two (arbitrarily chosen) transition

¹⁷Note that in contrast to Choo and Siow (2006), the probability that a worker of a given origin type o is chosen by a firm of destination type d depends on share of workers of type o in the population, $f(o)$. This difference stems from allowing an unobserved surplus component at the worker-position level. Menzel (2015) derives a similar formula in his nontransferable utility assignment model based on an unobserved surplus component at the agent pair level.

¹⁸Note that in the empirical work each transition will have both an origin and a destination because we do not observe unfilled vacancies and we augment \mathcal{K} to include a sufficient number of nonemployment “positions”. As a result, the number of workers seeking jobs will equal the number of destination “positions”.

group types g_1 and g_2 :

$$\begin{aligned}
\ln\left(\frac{\hat{P}_{g_1|d_1}}{\hat{P}_{g_2|d_1}}\right) &= \ln\left(\frac{e^{\frac{\theta_{g_1}}{\sigma}} P(z(g_1)|o(g_1), d_1) f(o(g_1)) C_{o(g_1)}}{\sum_{o' \in \mathcal{O}} \sum_{g' \in (o, d_1)} e^{\frac{\theta_{g'}}{\sigma}} P(z(g')|o', d_1) f(o') C_{o'}}\right) \\
&\quad - \ln\left(\frac{e^{\frac{\theta_{g_2}}{\sigma}} P(z(g_2)|o(g_2), d_1) f(o(g_2)) C_{o(g_2)}}{\sum_{o' \in \mathcal{O}} \sum_{g' \in (o, d_1)} e^{\frac{\theta_{g'}}{\sigma}} P(z(g')|o', d_1) f(o') C_{o'}}\right) \\
&= \frac{\theta_{g_1}}{\sigma} + \ln(P(z(g_1)|o(g_1), d_1)) + \ln(f(o(g_1))) + \ln(C_{o(g_1)}) - \\
&\quad \frac{\theta_{g_2}}{\sigma} - \ln(P(z(g_2)|o(g_2), d_1)) - \ln(f(o(g_2))) - \ln(C_{o(g_2)}) \tag{22}
\end{aligned}$$

Since the set $\{P(z(g)|o, d)\}$ is either observed or directly estimable (depending on whether a sample or the entire population is used), for the purposes of establishing identification we can treat $\ln(P(z(g_1)|o(g_1), d_1))$ and $\ln(P(z(g_2)|o(g_2), d_1))$ as known. Similarly, $f(o(g_1))$ and $f(o(g_2))$ are observed/estimable. Bringing all these terms to the left side of equation (22), we obtain:

$$\ln\left(\frac{\hat{P}_{g_1|d_1}/(P(z(g_1)|o(g_1), d_1)f(o(g_1)))}{\hat{P}_{g_2|d_1}/(P(z(g_2)|o(g_2), d_1)f(o(g_2)))}\right) = \left(\frac{\theta_{g_1} - \theta_{g_2}}{\sigma}\right) + (\ln(C_{o(g_1)}) - \ln(C_{o(g_2)})) \tag{23}$$

We see that the adjusted log odds only identifies the relative mean (re-scaled) surplus values from transition groups g_1 and g_2 , $(\frac{\theta_{g_1} - \theta_{g_2}}{\sigma})$, in the case where both groups are associated with the same origin job type: $o(g_1) = o(g_2)$. Otherwise, the difference in surplus values is conflated with the log difference in mean exponentiated worker discounted utilities between the two origin job types $(\ln(C_{o(g_1)}) - \ln(C_{o(g_2)}))$.

However, now consider two additional transition groups g_3 and g_4 that are both associated with some destination type d_2 such that $o(g_3) = o(g_1)$ and $o(g_4) = o(g_2)$.¹⁹ The four groups $g_1 - g_4$ can be chosen to represent the two pairs of transition groups that would be created by the two ways to match a given pair of destination positions to a given pair of workers. If we augment equation (23) by instead taking the log of the *ratio* of the (appropriately re-scaled) odds of g_3 and g_4 (conditional on d_2) and the (appropriately re-scaled) odds of g_1 and g_2 (conditional on d_1), we obtain:

$$\begin{aligned}
&\ln\left(\left(\frac{\hat{P}_{g_3|d_2}/(P(z(g_3)|o(g_3), d_2)f(o(g_3)))}{\hat{P}_{g_4|d_2}/(P(z(g_4)|o(g_4), d_2)f(o(g_4)))}\right)\right) / \left(\frac{\hat{P}_{g_1|d_1}/(P(z(g_1)|o(g_1), d_1)f(o(g_1)))}{\hat{P}_{g_2|d_1}/(P(z(g_2)|o(g_2), d_1)f(o(g_2)))}\right)\right) \\
&= \left[\left(\frac{\theta_{g_3} - \theta_{g_4}}{\sigma}\right) + (\ln(C_{o(g_3)}) - \ln(C_{o(g_4)}))\right] - \left[\left(\frac{\theta_{g_1} - \theta_{g_2}}{\sigma}\right) + (\ln(C_{o(g_1)}) - \ln(C_{o(g_2)}))\right] \\
&= \left(\frac{\theta_{g_3} - \theta_{g_4} - (\theta_{g_1} - \theta_{g_2})}{\sigma}\right) \tag{24}
\end{aligned}$$

¹⁹ d_2 could be (but need not be) the same destination type as d_1 .

Thus, we see that the appropriate log odds ratio can identify the expected gain in mean scaled surplus values from a swap among partners from any two end-of-year job matches. Note that these difference-in-differences do not preserve information about the baseline welfare of either worker types or firm types: the mean discounted value of each worker type and each destination position type gets eliminated during the differencing/conditioning, respectively.

However, the set of difference-in-differences $\Theta^{D-in-D} \equiv \left\{ \frac{(\theta_g - \theta_{g'}) - (\theta_{g''} - \theta_{g'''})}{\sigma} \forall (g, g', g'', g''') : o(g) = o(g''), o(g') = o(g'''), d(g) = d(g'), d(g'') = d(g''') \right\}$ preserves the crucial information about the relative efficiency of different matchings that exists in the observed transition group frequencies.

Specifically, in the next subsection we show that identification of the set of surplus difference-in-differences is sufficient to generate the unique assignment in counterfactual simulations that forecast the aggregate distribution of transition types $P(g)$ for any arbitrary change in either the marginal distribution of worker origin match types $f(o)$ or the marginal distribution of destination position types $h(d)$ (or both). Furthermore, if more than one observed matching is available, then σ can potentially be (roughly) estimated as well, allowing for a proper welfare analysis that calculates the approximate mean utility and profit gain for each worker origin type and firm destination type, respectively, from any such shifts in labor supply or demand.

Because we do not have data on unfilled positions, from this point forward we focus attention only on the subset of positions $\tilde{\mathcal{K}} \in \mathcal{K}$ that match with workers at the end of the period. Note that the number of singles on both sides of the market are assumed to be observed when proving identification in Choo and Siow (2006) and Galichon and Salanié (2015).

While ignoring unfilled positions is a drawback of our empirical work, note that each subset of assignments within a stable matching must also be stable. Thus, stability among the matrix of observed year-to-year transitions between dominant jobs based on \mathcal{I} and $\tilde{\mathcal{K}}$ is a necessary condition for stability of the full market transition matrix defined by \mathcal{I} and \mathcal{K} . Therefore, the relationships between the transition surpluses $\{\pi_{ijk}\}$ that we recover would not be reversed if data were augmented with additional transitions and unmatched agents.

In principal, though, unfilled positions may put upward pressure on wages that affect the division of surplus between workers and positions, even if they do not affect the final assignment of workers to positions. Furthermore, unfilled positions might potentially become filled in the wake of the potential labor supply and demand shocks that I simulate, so that the incidence measured in the counterfactual simulations described below may be slightly distorted. For example, a tornado that eliminates a number of local jobs may depress local wage levels far enough for previously unfilled positions to become the most desirable options for some local workers. To formally rule out such scenarios, we must assume that no unfilled

position is ever the second-best option for any worker who takes a job in the destination period, both in our data and in our counterfactual simulations.²⁰ This assumption implies that the unfilled positions do not affect either the allocation of workers to positions nor the division of surplus among them. Violations of this assumption will likely lead us to overstate the magnitude of welfare gains from positive shocks and the losses from negative shocks, since some previously filled positions would choose to remain vacant when increased local labor demand increases the equilibrium payoffs positions must offer to workers, and some previously vacant positions might choose to hire a worker when decreased competition lowers the required worker payoffs.

3.2 Counterfactual Simulations

This subsection demonstrates how to predict the way in which a set of workers (initially matched with a set of positions) would be reallocated to a new set of positions, given a particular job matching technology (i.e. collection of production functions, utility functions, and search and recruiting cost functions). In our empirical work, such counterfactuals will involve altering the distribution of destination positions by introducing labor demand shocks of various forms.

We can characterize the set of workers to be reallocated using their distribution across origin match types, $f^{CF}(o)$. The “CF” superscript indicates that this could potentially be a counterfactual distribution (e.g. capturing a proposed influx of refugees). Similarly, the set of counterfactual positions to be filled can be represented by its type distribution $h^{CF}(d)$, and the prevailing technology can be denoted $\{\theta_g^{CF}\}$. The values $f^{CF}(o)$, $h^{CF}(d)$, and $\{\theta_g^{CF}\}$ are all treated as inputs that are either observed or constructed by the researcher/policymaker. The goal is to use these inputs to predict the equilibrium distribution of transitions across transition groups, $\{g = 1, \dots, G\}$, as captured by $P^{CF}(g)$.

Consider the case of a manufacturing plant considering relocation. The immediate change in the location of a set of manufacturing and management positions that would occur is known by a local development board, and the existing group mean surpluses $\{\theta_g^{CF}\}$ have been estimated; the board wishes to predict the extent to which the plant relocation will decrease the probability of nonemployment and more generally increase the job-related utility among existing workers/job seekers in the local area versus workers arriving from neighboring or distant locales (and perhaps the profits of other local firms relative to more distant firms).

I assume that the unique counterfactual assignment also satisfies Assumptions 1 and 2

²⁰Instead, positions who eventually settle for other workers represent the best outside option for each worker, in the following sense: for each worker, there exists a position that ends up being filled by another worker that could increase its payoff by hiring the chosen worker at a wage yielding a higher utility to the chosen worker than their utility at any wage any unfilled position would prefer to remaining unfilled. (Of course, the filled position cannot increase its payoff by outbidding the position that actually hires the chosen worker, since the observed matching is presumed to be stable).

above, which implies that $\sum_{i:g(i,k)=g} e^{-r_i^{CF}} \approx P^{CF}(g|o(g), d(g)) f^{CF}(o(g)) |\mathcal{I}| C_o^{CF}$ for sets of group- and type-specific constants $\{P^{CF}(z(g)|o(g), d(g))\}$ and $\{C_o^{CF}\}$. I also assume that the average share of workers of each origin type who are existing employees of a random position in each destination type, $\{P^{CF}(z(g)|o(g), d(g))\}$, is known, and treat it as an input. In particular, when $f^{CF}(o) = f^{y'}(o)$ and $h^{CF}(d) = h^{y'}(d)$ for some observed year y' , then the appropriate existing employee fractions can be obtained via $P^{CF}(z(g)|o(g), d(g)) = P^{y'}(z(g)|o(g), d(g)) \forall g$, which is observed. Maintaining the assumed Type 1 distribution for the idiosyncratic component of the match surplus ϵ_{ik} , the counterfactual conditional choice probability $P^{CF}(g|d)$ can be expressed as:

$$P^{CF}(g|d) = \frac{e^{\frac{\theta_g^{CF}}{\sigma}} P^{CF}(z(g)|o(g), d) f^{CF}(o(g)) C_o^{CF}}{\sum_{o' \in \mathcal{O}} \sum_{g' \in (o, d)} e^{\frac{\theta_{g'}^{CF}}{\sigma}} P^{CF}(z(g')|o'(g'), d) f^{CF}(o') C_{o'}^{CF}} \quad (25)$$

The origin type-specific mean worker discounted utility values $\{C_1^{CF} \dots C_O^{CF}\}$ are equilibrium objects that will be affected by the counterfactual changes in technology incorporated into $\{\theta_g^{CF}\}$ and the counterfactual changes in the composition of both supply and demand unknown incorporated into $f^{CF}(o)$ and $h^{CF}(d)$, and are thus unknown. As a result, each counterfactual conditional choice probability cannot be immediately constructed, and thus must be treated as a function of the vector of type-specific mean worker exponentiated discounted utility values. We assume that no position captured by $h^{CF}(d)$ chooses to remain vacant in any the equilibria we seek.²¹ This implies that only relative prices of alternative origin types matter in determining the equilibrium assignment, so that we can normalize $C_1^{CF} = 0$ and write $P^{CF}(g|d) \equiv P^{CF}(g|d)(\mathbf{C}^{CF})$. Note that the inability to observe unfilled vacancies prevents the determination of the changes in absolute levels of worker utility following labor supply and demand shocks.

If the counterfactual conditional choice probabilities $\{P^{CF}(g|d)\}$ represent the group-level aggregation of an equilibrium (stable) assignment, they must be consistent with type-level market clearing. In particular, the aggregate demand for each origin type o implied by the (appropriately weighted) sum of conditional choice probabilities may not exceed the supply, captured by $f^{CF}(o)$.

Furthermore, in Appendix A1 I show that if an extra, dummy ‘‘position’’ type is added with mass equal to the share of workers who will be left unmatched (knowable in advance given our assumption that the number of positions to be filled is perfectly inelastic with respect to the vector $\{C_1^{CF} \dots C_O^{CF}\}$), then the augmented ‘‘demand’’ (including ‘‘demand’’ from nonemployment) will in fact equal supply for each type o .²² These dummy nonem-

²¹As discussed above, unfilled positions may be ignored if they do not affect the discounted utility values $\mathbf{C}^{CF} \equiv \{C_1^{CF} \dots C_O^{CF}\}$ that clear the market.

²²We formally prove this result as Proposition A1 in the appendix. The intuition behind the proof is that the unique assignment could have been computed by aggregating the supply-side stability conditions $\mu_{ik} = 1$ iff $k \in \arg \max_{k \in \tilde{\mathcal{K}} \cup 0} \pi_{ik} - q_k$ instead. This can be motivated via an assignment mechanism

ployment positions represent a computational mechanism for appropriately incorporating the mean surpluses workers obtain from nonemployment, $\{\pi_{i0}^i\}$.

Since the origin-type distribution $f^{CF}(\ast)$ must sum to one, we obtain the following $O-1$ market clearing conditions:

$$\begin{aligned} \sum_{d \in \mathcal{D}} h^{CF}(d) \left(\sum_{g: o(g)=2} P^{CF}(g|d, \mathbf{C}^{CF}) \right) &= f^{CF}(2) \\ \vdots \\ \sum_{d \in \mathcal{D}} h^{CF}(d) \left(\sum_{g: o(g)=O} P^{CF}(g|d, \mathbf{C}^{CF}) \right) &= f^{CF}(O) \end{aligned} \quad (26)$$

These market clearing conditions represent a system of $O-1$ equations with $O-1$ unknowns, $\{C_2^{CF}, \dots, C_O^{CF}\}$. Given a solution to this system, one can then construct the counterfactual transition probability for any transition group via $P^{CF}(g) = \sum_d h^{CF}(d) P^{CF}(g|d, \mathbf{C}^{CF})$.

Galichon and Salanié (2015) and Decker et al. (2013) each show that the probability distribution over transition groups $P^{CF}(g)$ that satisfies the stability and market clearing conditions is unique. Since any solution to this system also satisfies the stability and market clearing conditions, it must be the unique aggregate counterfactual equilibrium assignment. This line of reasoning reveals that we can construct a unique counterfactual transition between labor market allocations given any marginal distributions of origin and destination job matches ($f^{CF}(\ast)$ and $h^{CF}(\ast)$) and any vector of mean group surplus values $\{\theta_g^{CF}\}$ (and any vector of existing employee shares $\{P^{CF}(z(g)|o, d)\}$).

A final issue remains to be discussed. In the counterfactual labor demand shock simulations below, we will treat the full set of group-level joint surpluses $\Theta \equiv \{\theta_g \forall g \in \mathcal{G}\}$ as known, despite the fact that our identification argument in section 3 suggests that only the set of surplus difference-in-differences Θ^{D-in-D} are identified. In Appendix A2, we prove the following proposition:

Proposition 1:

Define the set $\Theta^{D-in-D} \equiv \left\{ \frac{(\theta_g - \theta_{g'}) - (\theta_{g''} - \theta_{g'''})}{\sigma} \forall (g, g', g'', g''') : o(g) = o(g''), o(g') = o(g'''), d(g) = d(g'), d(g'') = d(g''') \right\}$. Given knowledge of Θ^{D-in-D} , a set $\tilde{\Theta} = \{\tilde{\theta}_g \forall g \in \mathcal{G}\}$ can be constructed such that the unique group level assignment $P^{CF}(g)$ that satisfies the system of excess demand equations (26) using $\theta_g^{CF} = \tilde{\theta}_g \forall g$ and arbitrary marginal PMFs for origin and destination types $f^{CF}(\ast)$ and $g^{CF}(\ast)$ will also satisfy the corresponding system

where workers bid for positions instead of the reverse. But under the assumption that all positions will be filled, “demand” for positions of each destination type d will exactly equal supply. We then show that the group-level assignment that satisfies the resulting system of excess-demand equations (and is thus the unique group-level assignment consistent with an individual-level stable matching) will also satisfy the corresponding system of equations for the worker-side of the market, provided that the appropriate “demand” from dummy nonemployment positions is added.

of excess demand equations using $\theta_g^{CF} = \theta_g \forall g \in \mathcal{G}$ and the same arbitrary PMFs $f^{CF}(\ast)$ and $g^{CF}(\ast)$. Furthermore, denote by $\tilde{\mathbf{C}}^{\mathbf{CF}} \equiv \{\tilde{C}_1^{CF}, \dots, \tilde{C}_O^{CF}\}$ the market-clearing utility values that clear the market using $\theta_g^{CF} = \tilde{\theta}_g$, and denote by $\mathbf{C}^{\mathbf{CF}} \equiv \{C_1^{CF}, \dots, C_O^{CF}\}$ the market-clearing utility values that clear the market using $\theta_g^{CF} = \theta_g$. Then $\tilde{\mathbf{C}}^{\mathbf{CF}}$ will satisfy $\tilde{C}_o^{CF} = C_o^{CF} e^{-\frac{\Delta_o}{\sigma}} \forall o \in \mathcal{O}$ for some set of origin type-specific constants $\{\Delta_o\}$ that is invariant to the choice of $f^{CF}(\ast)$ and $g^{CF}(\ast)$.

Essentially, the proposition states that the identified set of surplus difference-in-differences Θ^{D-in-D} contains sufficient information to generate the unique counterfactual group-level assignment $P^{CF}(g)$ that would be consistent with the corresponding true set of surpluses Θ . Furthermore, the vector of utility premia $\tilde{\mathbf{C}}^{\mathbf{CF}}$ that clears the market using the artificial surpluses $\tilde{\Theta}$ generated from the surplus difference-in-differences will always differ from the “true” premia $\mathbf{C}^{\mathbf{CF}}$ that clear the counterfactual market under Θ by the same o -type-specific constants regardless of the compositions of supply $f^{CF}(o)$ and $h^{CF}(d)$ used to define the counterfactual.

While absolute levels of counterfactual mean utility by origin type are never uniquely determined (even when Θ is fully known), the existence of the “bias” terms $\{\Delta_o\}$ in Proposition 1 indicates that the *relative* levels of utility among origin types in counterfactual allocations (including the true “counterfactual” that was observed) are not identified. One cannot infer the level of relative utility among workers at different skill levels who start the time period in different locations. This inability to determine the existing division of surplus among workers and firms for any origin and destination type combination, which does not appear in Choo and Siow (2006) or Galichon and Salanié (2015), stems from the assumptions necessary to accommodate the lack of data on unfilled vacancies.

However, because the “bias” terms $\{\Delta_o\}$ are constant across counterfactuals featuring different supply and demand compositions $f^{CF}(o)$ and $h^{CF}(d)$, the relative *differences* in origin-type mean utilities, $(\frac{(\bar{r}_o^{CF1} - \bar{r}_o^{CF2}) - (\bar{r}_{o'}^{CF1} - \bar{r}_{o'}^{CF2})}{\sigma})$ among two counterfactuals across origin types can be identified. Note that such a pair of counterfactuals might include one that features a stimulus package or natural disaster versus an otherwise identical counterfactual that does not. For certain counterfactual scenarios, one might plausibly restrict utility gains for a particular origin type reference group to be known, thereby allowing (scaled) utility gains or losses $\frac{\bar{r}_o^{CF1} - \bar{r}_o^{CF2}}{\sigma}$ for other origin-type groups to be identified. Below, I assume that a plant relocation generates zero aggregate utility gains for workers (since aggregate labor demand did not change), creating a natural reference group (the population-weighted mean across all origin types) for whom the change is assumed to be zero. Alternatively, I assume that the small, very local stimuli and natural disasters I consider generate zero utility gain or loss for workers on the opposite side of the country. Such restrictions allow the share of total welfare gains (or losses) from a labor demand (and/or supply) shock experienced by each origin worker type to be determined. Furthermore, because the model is symmetric

with respect to workers and positions, this also implies that the mean changes in destination firm-type discounted profits $\frac{\bar{q}_k^{CF1} - \bar{q}_k^{CF2}}{\sigma}$ can also be identified, so that the share of profit gains (or losses) can also be computed. Thus, when combined with the available data, the model permits a reasonably complete analysis of welfare incidence from labor supply and demand shocks.

3.3 Identifying σ

While the share of welfare gains or losses for workers (or firms) can be identified without knowledge of σ , σ is nonetheless a parameter of consider interest. Because both the money-metric utility function (1) and the profit function (3) we adopted above are linear in worker earnings, knowledge of σ would allow the estimated scaled utility gains $\frac{\bar{r}_o^{CF1} - \bar{r}_o^{CF2}}{\sigma}$ for origin worker types and scaled profit gains $\frac{\bar{q}_k^{CF1} - \bar{q}_k^{CF2}}{\sigma}$ to be re-scaled in dollar terms, making it easy to gauge whether the utility gains or losses from a given labor market shock are economically meaningful.

Recall that σ captures the relative contribution of the idiosyncratic worker-firm components $\epsilon_{ij(i)k}$ versus systematic group-level components θ_g to the overall variance in joint surpluses across all worker-position pairs. Intuitively, when one observes a given destination position type C choosing an origin worker type A much more often than origin worker type B , it could be because $\theta_{AC} \gg \theta_{BC}$ even though σ is moderately large, or because θ_{AC} is marginally larger than θ_{BC} but σ is tiny. If the former is true, clearing the market after a shift in destination firm composition could require large changes in the utility values that must be promised to workers to engender sufficient substitution across worker types to overcome strong tendencies for particular workers and positions to match. If the latter is true, very small utility changes would suffice. Thus, σ determines the elasticity of worker choices with respect to relative wages, and is rather important in determining the degree to which changes in labor demand composition cause substantial reallocation of utility across skill types and geographic areas.²³

As Galichon et al. (2017) have noted, σ is not identified from a single observed matching.

²³Note that the existence of a single σ parameter governing relative wage elasticities is driven by assuming additive separability of the mean choice-specific values and an i.i.d type 1 extreme value distribution for the vector of idiosyncratic surpluses for particular worker/destination position combinations. As shown by Menzel (2015), the identification argument and the counterfactual simulations are not sensitive to the assumption of a type 1 extreme value distribution per se. Instead, the key assumption is that the idiosyncratic worker/position surpluses are independently and identically distributed according across all alternative worker/position matches, which leads choices by one side of the market to satisfy the independence of irrelevant alternatives (IIA) property conditional on required values to be offered the other side. As the relevant matching market gets large (as the U.S. national labor market most certainly is), the counterfactual allocation associated with any joint distribution of idiosyncratic surpluses satisfying this property will converge to a common, unique allocation. Galichon and Salanié (2015) and Chiappori et al. (2009) discuss the possibility of relaxing the *i.i.d* assumption to allow certain forms of correlation in the idiosyncratic component across matches featuring shared characteristics. However, I maintain the standard i.i.d. assumption in this paper in order to minimize an already substantial computational burden.

However, combining information from multiple matchings can potentially identify σ . Since I observe national market-level matchings (transitions from a set of origin job matches to a set of destination job matches) for each pair of years 1993-1994 to 2010-2011, I attempt to calibrate a value of σ .

I exploit the fact that the composition of U.S. origin and destination job matches $f^y(o)$ and $h^y(d)$ evolved across years y . Specifically, I estimate the set of group-level surpluses $\{\theta_g^{2007}\}$ from the observed 2007-2008 matching. Then, holding these surplus values fixed, I combine $\{\theta_g^{2007}\}$ with $f^y(o)$ and $h^y(d)$ from each other year $y \in [1993, 2010]$ to generate counterfactual assignments and changes in mean (exponentiated) scaled utility values $\{C_o^{CF,y}\}$ for each origin type. These counterfactuals predict how mean worker utilities by skill/location combination could have been expected to evolve over the observed period given the observed compositional changes in labor supply and demand had the underlying surplus values Θ been constant and equal to Θ^{2007} throughout the period.

I then regress the actual mean annual earnings changes experienced by different origin types o from origin periods $y-1$ to destination periods y on logs of the predicted changes in mean scaled utility values $\{\ln(C_o^{CF,y})\}$, which approximately equal $\frac{\bar{r}_o^{CF,y}}{\sigma^y}$ if individual-level utility changes r_i^{CF} are roughly similar within origin type o . To the extent that a) most of evolution in the utility premia enjoyed by workers in particular locations and skill categories was due primarily to changes in supply and demand composition rather than changes in the moving costs, recruiting costs, tastes, and relative productivities that compose the joint surplus values Θ , and b) mean utility gains for each origin type generally came from increases in mean annual earnings in the chosen year rather than increases in non-wage amenities or continuation values, the coefficient on $\frac{\bar{r}_o^{CF,y}}{\sigma^y}$ will approximately equal σ^y . Appendix A3 provides further detail on this procedure.

Clearly, given the additional strong assumptions required, this approach represents a relatively crude attempt to calibrate σ . In practice, the estimates we obtain for σ are surprisingly consistent across years. We used the mean estimate of σ^y across all sample years, $\bar{\sigma} = 8,430$, to produce dollar values for all the results relating to utility gains presented below.

As noted by Galichon and Salanié (2015), when match-level unobserved heterogeneity is assumed away as in Choo and Siow (2006), the observed earnings of the matches between origin workers and destination positions can be further used to attempt to decompose each group-level joint surplus value θ_g into its worker and firm subcomponents (denoted θ_g^l and θ_g^f respectively). In other words, one can determine the relative contributions of amenities/future earnings opportunities versus current and future revenue contributions to θ_g . However, in Appendix A4 I show that clean identification of θ_g^l and θ_g^f breaks down without the particular structure Choo and Siow (2006) place on the unobserved match component ϵ_{ik} unless further strong assumptions are imposed. I chose not to pursue this path further in this paper, primarily because I have shown that this decomposition is

unnecessary to determine the incidence across worker and firm types of alternative local labor demand shocks, which is the primary goal of the paper.

3.4 Interpreting the Counterfactual Simulations

The rest of the paper exploits the insights from the assignment model to conduct an analysis of geographic labor market integration. The analysis involves constructing counterfactual assignments that illustrate how a local development agency might attempt to forecast the impact of alternative local labor demand shocks on the welfare of local (and less local) residents. While the main focus for these simulations is the heterogeneity across alternative shock compositions, future drafts will also examine how the payoff to local development policy might have changed over time by performing this calculation using mean group surplus values from different generations.

However, there are several important caveats relating to the theoretical model as it applies in the spatial labor market context that merit mentioning before proceeding to the details of data and estimation.

First, when simulating shocks to the level and composition of labor demand below, I will generally treat as fixed the job matching technology defined by the set of identified joint surplus difference-in-differences, Θ^{D-in-D} . However, for the simulated allocation and incidence measures to be accurate, this implicitly assumes that relocating, adding, or subtracting positions of different types from or to particular local areas does not itself change the determinants of these joint surplus values. These include worker preferences for locations, the revenue that different types of workers generate for different types of firms, the search, moving, and recruiting costs associated with changing positions, and even the continuation values associated with beginning the following period with a particular position or a particular worker.

There are a couple of particularly plausible violations of this assumption. First, to the extent that production agglomeration economies exist, the existence of a new firm nearby might increase the demand for intermediate products produced by other local firms, thereby raising the productivity of workers for such firms. Second, if the new jobs are thought to be persistent and search/recruiting/moving costs increase with distance, then obtaining a job at a different firm in the same local area as a newly relocated firm might now have greater continuation value because future job searches will begin in a local area featuring a higher level of labor demand.

However, while both of these scenarios change the joint surpluses θ_g of transition groups involving nearby firms, if the increased productivity of workers or the increased continuation value for workers at such firms is common to all potential origin types o , then it will not affect the surplus *difference-in-differences* Θ^{D-in-D} that generate the counterfactual assignment. Furthermore, the proof of Proposition 1 in Appendix A2 shows that such composition-

induced surplus changes also will not affect the equilibrium origin type payoffs that capture the shock’s incidence among groups of workers.²⁴ Instead, any increase in joint surplus by a destination position k that is common to all workers will be fully reflected in k ’s profit payoff, either through higher revenue for the same costs (agglomeration case) or through lower salaries that offset the change in worker continuation value.²⁵ Thus, bias in forecasted worker incidence in counterfactuals from shock-induced changes in joint surpluses only stems from differential changes in joint surplus among origin worker types for a given destination type.

Along the same lines, we also must assume that the labor demand shock does not induce further changes in either firms’ location decisions or the number of positions they wish to fill. To the extent that firm relocations or startups cause other firms to form, relocate, or expand, these additional compositional changes would need to be anticipated and incorporated into the simulated shock to capture the net change in destination type composition, perhaps based on estimates from other spatial models that focus on product markets.

A second caveat relates to the permanence of the shock. For our stimulus packages we generally assume that the new positions generate the same surplus values θ_g as existing positions of the chosen destination type. Implicitly, this requires that they have the same expected duration over time as any other position of their type. If one wished to simulate a temporary construction stimulus, one would need to estimate a separate set of surplus parameters for temporary versus permanent construction jobs. More generally, a more precise distinction of differences in welfare effects between shocks of different expected durations requires a fully dynamic assignment model along the lines of Choo (2015). Similarly, the Markov-style model used here conditions only on a single previous period’s location, and thus cannot accommodate optimal sequences of location decisions and return migration choices highlighted by Kennan and Walker (2011) that would enrich a dynamic analysis.

Note that in principle, one could simulate how a given shock changes the sequence of year-to-year assignments over a longer horizon. This would involve using the simulated assignment after one year to update the distribution of origin types for the following year, then simulating a second year-to-year aggregate matching using the updated origin type distribution (and perhaps an updated destination type distribution, if one predicted additional firm entry caused by the shock, as discussed above). Continuing this simulation process for many years would yield an ergodic, steady-state joint distribution of job match observable characteristics. One could then compare this distribution to the ergodic distribution that would emerge in the absence of the shock.²⁶ We do not pursue this approach in this paper,

²⁴Using the proof’s notation, such surplus changes will only change Δ_d^2 , which does not enter into equilibrium mean payoffs for origin types $\{C_o^{CF}\}$.

²⁵Note, though, that in these scenarios the profit gains among nearby firm types d will be understated. The possibility of differential agglomeration effects for nearby firms across different shock compositions (emphasized by Glaeser et al. (2008)) is one reason that we focus primarily on incidence among workers, for whom differential agglomeration effects are likely to be of second order importance.

²⁶CDP adopt this approach in the context of their trade-centric dynamic spatial equilibrium model.

in part because appropriate updating of the shares of workers in each origin earnings quintile requires accurate predicted earnings gains, which relies heavily on the accuracy of the estimate of σ , which is likely to feature some bias.

A final, important caveat relates to the absence of a housing market in the model (and the corresponding absence of residential choices in the data). Standard models of spatial equilibrium in urban economics (e.g. Roback (1982) or Kline and Moretti (2013)) emphasize the critical role of the housing market in determining incidence from place-based policies. In particular, if housing supply is perfectly inelastic and workers are sufficiently mobile, the entire incidence of a positive place-based shock is enjoyed by landholders in the form of higher rents (which fully offset the utility gains to workers from any wage increases). Thus, in principle failing to model the housing market could result in highly biased estimates of shock incidence.

However, in contexts where housing supply is likely to be relatively elastic (such as rural areas or areas with weak zoning laws) or where there exists excess housing supply due to a declining population, housing prices may move little, and abstracting attention from the housing market may produce little bias in incidence forecasts. Indeed, Gregory (2013) finds that neighborhoods receiving empowerment zone status, a local labor demand shock similar to those we estimate, experienced negligible changes in rent but substantial wage gains among residents, suggesting that omitting a housing market response might generate minimal bias.²⁷

Along the same lines, in contexts where tastes for particular neighborhoods is strong but commuting between neighborhoods is fairly low cost, much of the adjustment to shocks may take the form of changing commuting patterns, with very little change in demand for housing across locations. Indeed, since commuting costs from job transitions that involve locational changes already constitute a component of the joint surplus θ_g , they are appropriately captured by the model; thus a scenario in which the true response to a labor demand shock involves lots of small commuting adjustments that ripple out from the focal point of the shock is likely to be closely matched by the model-based simulations, with potentially accurate incidence estimates.²⁸ While a full welfare analysis certainly requires incorporating housing and product markets, the goal of the approach taken here is to highlight the heretofore underappreciated role of differential geographic scopes of local labor markets for different types of workers and firms in determining the incidence of alternative local labor demand interventions.

²⁷The authors note that the neighborhoods receiving these shocks were often in locations experiencing recent decline that were unattractive residential options for many, so that the lack of impact on local rent rates may not generalize to shocks to healthier locations.

²⁸Note also that mobility frictions induced by housing markets are likely to be partly reflected in the log odds ratios capturing the relative propensities with which different origin worker types make certain types of job transitions that are used to identify the set Θ^{D-in-D} . So differential willingness to pay high prices for locational amenities will be captured by heterogeneity in $\{\theta_g\}$ across origin worker types for groups involving positions in the same location, and thus will be reflected in the counterfactual simulations.

4 Data

4.1 Overview

I construct a dataset of year-to-year worker job transitions (pairs of primary jobs in consecutive years) using the Longitudinal Employer-Household Dynamics (LEHD) database. The core of the LEHD consists of state-level wage records collected for unemployment insurance purposes that contain quarterly worker earnings and unique worker and firm IDs for a near universe of jobs in the state.²⁹ The worker and firm IDs are then linked across states, and the data are augmented with information on firm- and establishment-level characteristics (notably establishment locations and firm-level detailed industry codes) from a state-supplied extract of the ES-202/QCEW report and individual-level data from the Social Security Administration (including age, race and sex but not including occupation nor education for most of the sample).³⁰

4.2 Sample Selection

I restrict attention in this preliminary draft to a sample of 17 states that provide data to the LEHD system as of 1993: AZ, CA, CO, FL, ID, IL, IN, KS, LA, MO, MT, NC, OR, PA, WA, WI, WY.³¹ This sample was chosen so as to be able to illustrate how the geographic scope of labor markets has changed over the last generations. While this draft focuses on simulations based on only 2010-2011 worker transitions between primary jobs, future versions of the paper will examine changes in the projected geographic incidence of shocks over time and will consider an augmented set of states. We also restrict the sample to person-years featuring individuals with ages between 20 and 70. This restriction limits the influence of “nonemployment” spells consisting of full-time education or retirement followed by part-time work, so that parameters governing nonemployed workers would be identified primarily from prime-aged workers who were unemployed or temporarily out of the labor force.³²

The resulting dataset is then converted from a job-quarter-year-level dataset to one whose observation level is the combination of a person and a pair of primary jobs in consecutive years (i.e. person-level job transition or retention). This is done by first identifying each individual’s primary job in each year, then aggregating earnings from the primary job

²⁹The database does not include farm jobs or self-employed workers. I also exclude federal employees, who must be merged in via a separate OMB database. This has little consequence given that the current sample does not contain Virginia, Maryland, and the District of Columbia.

³⁰For further details about the contents and construction of the LEHD, see Abowd et al. (2009).

³¹Alaska and Maryland also provide data as of 1993, but we exclude Alaska due to its geographic isolation and Maryland due to difficulty in merging in records of federal employees (Maryland workers who became federal employees would show up as non-employed).

³²For computational reasons, the results in this draft are based on a 50% random sample of all transition-level observations from the sample just defined.

across all quarters within the year, and then appending primary jobs from the following year to the current observation to create a transition/retention observation. The primary job for a worker is defined as the job with the highest earnings that exists for at least one full quarter (a job is observed in a full quarter if the worker-firm pair reports positive earnings in the preceding and following quarter as well). A worker who does not report earnings above \$2,000 at any job in any full-quarter in a given year in any observed state (including observed states outside of the chosen sample of states for the more recent years) is designated nonemployed. An individual is included in the sample if he/she is ever observed as employed in one of the sample states between 1993 and 2010. I remove from the estimation sample observations associated with worker year-pairs in which the worker is observed working in another non-sample state in either year $y - 1$ or y in order to prevent spurious transitions in and out of nonemployment where possible. This amounts to treating such workers as if they were out of the relevant labor force for the sample states in the chosen years. (The results in this preliminary draft focus on the 2010-2011 transition in which all states but Massachusetts are observed, which limits the problem of spurious transitions into and out of nonemployment driven by employment transitions into and out of the set of observed states).³³

I limit each sample member's years in the sample to consist of years between his/her first and last years of observed employment. Thus, the sample only includes spells of nonemployment that are bookended by spells of observed employment. I exclude nonemployment spells before the first year and after the last year of observed employment for multiple reasons. First, I hope to minimize the incidence of spurious nonemployment in which the worker is working out of the chosen sample of states. Second, we want our counterfactual simulations measuring the incidence of shocks on existing non-employed workers to be based on the relative tendencies of alternative firm types to hire nonemployed workers rather than first time entrants, so that the interpretation of such incidence measures is straightforward. I suspect that job applications of first-time entrants and prime-age unemployed workers are treated very differently by employers. The next version of the paper will reincorporate new entrants at each location into the model as separate origin types o .

The removal of nonemployment spells at the beginning and end of each worker's employment history creates a bias toward including too few spells of nonemployment near the beginning and end of the sample, because in the first (last) year any nonemployed worker could not possibly have been observed as working in a prior (next) year. I address this by imputing the missing nonemployment-to-employment, employment-to-nonemployment, and nonemployment-to-nonemployment spells using the pattern of employment-to-employment

³³Since the results in this draft were released, I have lost access to the full set of transitions from all the states. Thus, the next draft will solve the problem of spurious transitions to and from "nonemployment" driven by moves to and from jobs in out-of-sample states by instead merging public-use ACS data capturing employment-to-employment transitions by demographic group for all 50 states. The ACS data will also allow me to distinguish unemployed workers from workers who are out of the labor force.

transitions during the sample (which our sample selection procedure does not affect) combined with the relationship between the distribution of E-to-E transitions and NE-to-E, E-to-NE, and NE-to-NE transitions, respectively, during the middle years of our sample (when this sample selection bias should be minimal). The details of this imputation procedure are presented in Appendix A6.³⁴

However, because I exclude job-to-job transitions into or from states outside of our 17 state sample, the counterfactual simulations will likely overstate the geographic concentration of demand shock incidence, since workers from the remaining states are effectively excluded from competing for the new positions.³⁵

4.3 Assigning Job Matches to Types and Job Transitions/Retentions to Groups

For each pair of years $(y - 1, y) \in \{(1993, 1994), \dots, (2010, 2011)\}$ we assign each job transition/retention observation to an origin type $o^{y-1}(i, j)$, and destination type $d^y(k)$, and a transition group $g^{y-1,y}(i, j, k)$ (time superscripts will henceforth be dropped except where necessary). Specifically, a worker i with primary year $y - 1$ position j is assigned to an origin type $o(i, j)$ based on the combination of the location of the establishment associated with position j ($m(j)$) and the earnings quintile at primary position j .³⁶ The same worker i with year y primary position k is assigned to a destination type $d(k)$ based on the combination of the geographic location associated with establishment $m(k)$, its size quartile (based the establishment employment distribution), its quartile of average worker earnings (again using an establishment-level distribution), and the industry supersector of the firm associated with the establishment. The pair of primary jobs for worker i , (i, j, k) , is assigned to a group $g(i, j, k) \equiv g(o(i, j), d(k), z(i, j, k))$ based on the origin job type $o(i, j)$, the destination job type $d(k)$, and an indicator $z(i, j, k)$ for whether establishments for positions j and k are the same ($m(j) = m(k)$).

³⁴In the next draft I will exploit ACS and CPS data to generate counts of nonemployment-to-employment, employment-to-nonemployment, and nonemployment-to-nonemployment transitions by transition group g rather than relying on the absence of employment in the LEHD, so this imputation procedure will not be necessary.

³⁵In the next draft, I will examine sensitivity of the incidence estimates to excluding nearby states by considering results from a subset of simulations featuring focal tracts in states where we observe all or nearly all of the surrounding states. Results from these simulations exist but have not yet been approved for release to the public. However, all of the qualitative results discussed below continue to hold. In particular, the new simulations show a similarly large share of the welfare gains accrue to out-of-state workers.

³⁶Earnings quintile cutoffs are defined relative to the distribution of primary job annual earnings for workers in the state-year combination associated with the observation, and are based on prorating earnings in full quarters only to ensure that the quintile better captures a workers' skill rather than the share of the year in which they were working. One drawback of the LEHD database is that a worker's location must be imputed for multi-establishment firms. However, the Census Bureau's unit-to-worker imputation procedure assigns an establishment to a worker with a probability that decreases in the distance between the worker's residence and that establishment. Consequently, cases of significant measurement error in true location are unlikely to occur, since most mistakes will misattribute the worker's job to an establishment within the same tract or perhaps a nearby tract

5 Estimation

5.1 Defining the Local Labor Demand Shocks

I consider three categories of local labor demand shocks: stimulus packages, firm relocations, and natural disasters. Each stimulus package shock consists of 500 jobs that are added to the destination-year stock of jobs to be filled in a chosen census tract, combined with the removal of 500 nonemployment “positions”. Given that census tracts have on average around 5,000 jobs, this represents about a 10% increase in labor demand for the average tract. For each chosen tract, I simulate 55 different stimulus packages, each differing in the particular kind of firm whose demand increases by 500 jobs based on combinations of the remaining non-location firm attributes that defined a destination type in the model above: firm size quartile, quartile of average worker earnings, and industry supersector. The full list of the alternative firm compositions of the simulated shocks is displayed in Table 1. The compositions were chosen to highlight the heterogeneity in incidence across different industry/firm size/firm earnings category cells.

The firm relocation shocks are nearly identical in structure to the stimulus packages. Each relocation shock contains a 500 job increase assigned to a particular census tract, and the same 55 shock compositions are considered as for the stimulus packages. The only difference is that instead of subtracting 500 positions from the nonemployment destination type, the 500 positions are instead subtracted from a sending destination type associated with the same firm characteristics as the type receiving the jobs but located in a far away state (at least two states away, so that locations near to the “winning” site are minimally affected by the lost employment at the “losing” site). Unlike the stimulus package shocks, which increase nationwide labor demand, the relocation packages merely redistribute existing demand across locations. If there were no spatial search frictions and workers could move costlessly, then such shocks would merely cause the moving firm’s workers to follow the firm, no additional reallocation would be necessary, and the “shock” would have zero incidence for workers in all initial skill groups and locations. By contrast, if mobility costs were so high as to eliminate mobility between small local labor markets, the impact on local workers of a stimulus package and a relocation with the same firm composition would be identical.

Finally, I also consider “natural disaster” shocks in which a targeted census tract loses a random 25%, a random 50%, or all 100% of its jobs in the destination year, with the number of lost jobs being added as “positions” to the nonemployment destination type. These simulations provide an opportunity to examine whether the skill incidence of negative shocks is symmetric to positive shocks, as well as to consider the degree to which higher skilled workers initially working in the targeted tract are able to capture a disproportionate share of the remaining local jobs when only a fraction of such jobs are eliminated. These disaster

simulations are also included to illustrate how the two-sided matching model could be customized to handle any particular disaster scenario, including disasters such as hurricanes that hit a number of contiguous tracts simultaneously (and perhaps with differential force).

To this point, I have used a census tract as the unit of aggregation at which to define a geographic location. However, since 1) origin types consist of tract \times Earnings Quintile/Nonemployment cells, 2) destination types consist of tract \times firm size quartile \times firm avg. earnings quartile \times industry supersector cells, 3) origin-destination combination featuring the same tract can have both job stayers and local job switchers, and 4) the chosen 17 states contain 32,837 census tracts, there are $32,837 \times 6 \times 32,837 \times 4 \times 4 \times 11 + 32,837 \times 6 \times 4 \times 4 \times 11 = 1.13 \times 10^{12}$ groups in the initial group space. Even with upwards of 25 million transitions for a given pair of adjacent years populating the estimates $P(\hat{g})$, I am still left with only 0.00002 observations per element of the group-level distribution, so that the empirical group distribution $P(\hat{g})$ is too noisy to use directly in place of the underlying distribution $P(g)$ that might be observed with a much larger population. I address this problem in two ways.

5.2 Collapsing the Type Space for Distant Geographic Areas

First, since I am particularly interested in the incidence of demand shocks of alternative compositions across locations relatively near to the site of the shock, I combine groups that are defined by the same worker and firm characteristics and are geographically proximate to each other but far from the site of the shock. Specifically, given the focus on carefully measuring local incidence, I do not combine any origin or destination types featuring a location within a five tract pathlength of the targeted tract for a given simulation. However, I combine into a single origin (or destination) type any origin (or destination) types that share the same worker (or firm) characteristics and feature tracts that are in the same state as the targeted tract and in the same public-use microdata area (PUMA) as each other (but not the same PUMA as the targeted tract). Thus, outside a 5-tract circle surrounding the targeted tract, the geographic locations of types are defined by PUMAs rather than tracts. Furthermore, for types featuring tracts outside the targeted state, we combine types featuring the same worker (or firm) characteristics whose tracts are in the same state. Thus, outside of the targeted state, the geographic locations of types are defined by states rather than PUMAs.

Coarsening the type space for distant geographic locations dramatically decreases the overall number of groups and the severity of the sparse matrix problem. In particular, while many job-to-job transitions are between nearby tracts, there are very few transitions between any chosen tract in California and a corresponding tract in Kansas, so that relative surplus parameters for transition groups featuring tracts in different states would never be well-identified without such coarsening. Note that this approach still incorporates all of the

transitions and all of the locations in the 17 state sample into each simulation, so that each local labor market is still nested within a single national labor market.

However, this type aggregation procedure does imply that the origin and destination type space will be different for each simulation that involves a different targeted tract. Furthermore, disclosure restrictions imposed by the FSRDC system prevent the release of any results that are specific to a particular substate geographic location. Thus, while each simulation is performed with a particular tract level target, I only report averages of incidence measures across 500 different simulations for each shock type, where each simulation features a different randomly chosen target census tract from the sample.³⁷ Importantly, the same set of 500 randomly chosen target census tracts is used for each of the alternative simulations, so as to facilitate fair comparisons between the alternative stimulus packages, firm relocations, and natural disasters.

After the simulations have been run, the space of transition groups is again redefined in order to average simulation results across alternative targeted census tracts. This time, origin and destination type locations are replaced with bins capturing distance to the targeted census tract, and I report estimates of incidence for various distance rings around the site of the shock. Note, though, that during the simulations themselves the spatial linkages between adjacent and nearby tracts are not restricted to follow a particular parametric function of distance between locations. Thus, to this point no prior assumption about the role of distance has been imposed during estimation.

5.3 Smoothing the Empirical Group-Level Distribution $P(\hat{g})$

The second approach used to overcome the sparsity of the empirical group-level distribution $P(\hat{g})$ involves smoothing this distribution prior to estimation by making each element’s value a kernel-density weighted average of groups featuring “similar” worker and firm characteristics. Such smoothing introduces two additional challenges. First, while such weighted averages increase the effective sample size used to estimate each element of $P(\hat{g})$, excessive smoothing across other transition groups erodes the signal contained in the data about the degree of heterogeneity in the relative surplus from job transitions featuring different combinations of worker characteristics, firm characteristics, and origin and destination locations. Since highlighting the role of such heterogeneity is a primary goal of the paper, decisions about the appropriate smoothing procedure must be made with considerable thought. A second (but related) challenge consists of identifying which of the worker and firm characteristics that defines other cells makes them “similar”, in the sense that the surplus $\theta_{g'}$ of an alternative group g' that shares particular observed characteristics with group g is likely to closely approximate the surplus θ_g whose estimate we wish to make more precise.

³⁷I do impose that a census tract is only eligible to be a target tract in the simulations if it features at least 100 jobs, so that the parameters governing the behavior of local firms and workers is reasonably well-identified.

I base the smoothing procedure on the intuition that the geographic location of a destination establishment is likely to be critical in determining the origin locations whose associated worker transitions generate the most surplus (i.e. least moving/search cost), while the combination of non-location characteristics (firm size, firm average worker earnings, and firm industry) is likely to be more important than location in determining the skill category of worker (proxied by initial earnings quintile) that generates the most surplus.

Specifically, I use the fact that $P(g)$ can be initially decomposed into a conditional choice probability and a marginal destination type probability, $P(g|d(g))h(d(g))$, and further decomposed into $P(loc(o(g))|earn(o(g), z(g), d)P(earn(o(g), z(g)|d)h(d(g))$. I then construct separate kernel density estimators for $P(loc(o(g))|earn(o(g), z(g), d)$ and for $P(earn(o(g), z(g)|d)$.

The first term captures the probability that a hired worker’s initial location would be $loc(o(g))$, conditional on the prior earnings/nonemployment of the worker, whether the worker is a new or retained employee, and the destination type of the firm (including the firm’s location). The estimator for this component only smooths among alternative groups g' featuring the same origin and destination location pair as g , and gives greater weight to alternative groups featuring similar values of the other worker and firm characteristics that define a group (worker earnings and firm size, average earnings, and industry).³⁸

The second term captures the probability that a position of type $d(g)$ would be filled by a worker featuring a particular combination of prior earnings category and existing vs. new employee status (irrespective of worker location). The estimator of this term only smooths among alternative groups g' featuring the same combination of firm size, average earnings, and industry categories, and gives greater weight to groups whose destination location $loc(d(g'))$ is closer to the chosen groups destination location $loc(d(g))$. Further detail about the smoothing procedure is provided in Appendix A5.

This customized smoothing procedure has a number of desirable properties. First, by requiring the same origin and destination locations as a necessary condition for non-zero weight when estimating the propensity for particular destination types to hire workers from each location, it can generate considerable precision in estimated conditional choice probabilities without imposing any assumption about the spatial links between locations. Second, at the same time, it can still use information contained in the hiring and retention choices of more distant firms to learn about the propensity for firms of different size, pay level, and industry to hire workers at different skill levels and from nonemployment. Third, the procedure places non-trivial weight on transition groups featuring less similar worker and firm characteristics only when there are too few observed hires/retentions made by firms

³⁸There are a very small number of destination or origin types are never observed in any transition. By necessity, we put positive weight on groups featuring nearby origin or destination locations in such cases. However, the surpluses associated with these groups have no impact on counterfactuals unless the simulation generates new positions of the never-observed destination types or new origin workers of the never-observed origin types. We choose our target census tracts in the simulations below to contain sufficient observed employment to prevent this scenario.

associated with groups featuring very similar characteristics to yield reliable estimates.³⁹

5.4 Standard Errors

While this draft uses a 50% subsample of all pairs of dominant jobs in adjacent years in the population of 17 states we consider, the next draft will use the full 100%. Thus, it is not obvious how to define the relevant population for the purposes of inference. Furthermore, since I estimate millions of surplus parameters $\theta_g \in \Theta$, and each counterfactual incidence statistic depends on the full set Θ , I do not report estimates of any single parameter. Instead, any standard errors that might be reported ought to provide information about the precision of the incidence forecasts. Given that each incidence statistics provided in the next section is based on averaging across 500 simulations featuring different randomly chosen census tracts, a natural indicator of precision might be the standard deviation in the chosen incidence statistic across the 500 tract-specific simulations. These standard deviations have not yet passed through disclosure review, but will be available in the next draft (no standard errors appear in the tables below). However, while these standard deviations capture the variation in estimated θ_g parameters across geographic locations, they do not provide information about the degree to which this variation is caused by sampling error (relative to other possible U.S. labor markets drawn from the same data generating process) versus systematic differences in worker and firm productivities, moving costs, and non-wage utilities driven by past sorting. I will continue to examine other alternatives.

6 Results

6.1 Stimulus Packages

6.1.1 Incidence by Distance to Focal Tract

Before comparing the impacts of stimulus packages featuring different firm compositions, I focus first on characterizing the geographic scope of labor markets for a “typical” local stimulus. I do this by averaging the predicted change in worker-position allocations produced by the stimulus packages across all 32 stimuli simulated, effectively integrating over the distribution of industry, firm size, and firm average pay composition. While I focus attention on graphical representations of the results contained in a set of figures, most figures have an accompanying Appendix table (listed in parentheses) that contains the particular values that were plotted and are cited in the text.

³⁹In addition, the estimate of each element of the smoothed group-level distribution $P(g)$ places non-zero weight on many different groups, so that no element of the smoothed group-level distribution that results contains identifying information about any particular worker or firm, eliminating any disclosure risk (a necessary condition for releasing results out of the restricted FSRDC environment).

Figure 1 (Table A1, Col. 1) illustrates the mean probability of receiving one of the 500 new stimulus jobs for randomly chosen individuals initially working at different distances from the census tract receiving the stimulus. The figure reveals the sense in which U.S. labor markets are still extremely local: the probability of obtaining one of the new jobs for a worker initially working within the target tract (.032) is more than 6 times higher than for a worker working in an adjacent tract (.005), 13 times higher than for a worker working 2 tracts away (.002), and 30 times higher than for a worker working 3 or more tracts away with the same PUMA. Furthermore, additional distance from the focal tract continues to matter at greater distances: the probability of obtaining one of the stimulus jobs for a local (target tract) worker is 55 times higher than for a worker in an adjacent PUMA, 123 and 241 times higher than for a worker two PUMAs away or 3 or more PUMAs away within the same state, respectively, and 1,181 and 24,378 times higher than for a random worker one or two states away or 3 or more states away, respectively.

However, the vast differences in $P(\text{new job} \mid \text{distance from target})$ present a very misleading guide to the overall incidence of new jobs across geographic locations. This is because the target tract initially employs an extremely small fraction of the U.S. population defined by the model to be at risk of obtaining the stimulus jobs. Figure 2 shows the share of the workers in the simulation samples that are working in each distance bin relative to the targeted census tract prior to the stimulus. Only 0.0039% of the workforce is composed of workers initially working in the target tract. As expected, the shares get larger quickly as we move toward distance bins defined by concentric circles with much larger radii: 0.02, 0.05, 0.2 percent of the workforce initially work 1, 2, or 3+ tracts away from the target tract within the same PUMA, while 0.3%, 0.9% and 1.3% initially work 1, 2, or 3+ PUMAs away, respectively, and 16% and 81% initially work 1-2 states or 3+ states away.

Consequently, if one swaps the terms in the conditional probability and calculate the share of stimulus jobs obtained by workers initially working or nonemployed in each of the distance bins listed above, $P(\text{distance from target} \mid \text{new job})$, one obtains a very different impression of incidence. Figure 3 (Table A2, Col. 1) displays the mean share of new jobs by distance bin across the 32 simulated stimulus packages. 7.5% of new jobs go to workers initially in the target tract, another 28% are obtained by other workers in the PUMA, 33% are obtained by workers in different PUMAs within the state, and 31% are obtained by out of state workers. So a very large share of the new jobs are likely to be obtained by workers far outside the local jurisdiction that is hosting the stimulus (and is likely lobbying for its local placement).

One could likely obtain similar forecasts of the shares of workers by distance bin who would obtain jobs at a new firm simply by looking at the distance composition of workers who obtained jobs from actual plant openings in the past. As emphasized in the introduction, though, the probabilities of obtaining the particular new jobs created by the stimulus package may not be particularly informative about the true incidence of the shock. This is

because many of the workers who obtain the new jobs would have obtained other similarly paying jobs in the absence of the stimulus, and nearby workers may now obtain the jobs these workers would have accepted or retained, and so on, creating ripple effects through vacancy chains that determine the true employment and wage incidence. This is where the use of a flexible equilibrium model is particularly valuable.

Figure 4 (Table A3, Col. 1) is analogous to Figure 1 (Table A1, Col. 1), except that instead of the probability of obtaining a particular stimulus job, it captures the change in the probability of any employment (or equivalently, the change in the probability of nonemployment) due to the stimulus, relative to a no-stimulus counterfactual, for randomly chosen workers initially working at different distances from the target census tract of the stimulus.

The figure demonstrates that the change in employment probability is still quite locally concentrated, though less so than the probability of obtaining a stimulus job. Workers initially employed (or nonemployed) in the target tract are 1.2% more likely to be employed at the end of the year than in the absence of the stimulus. This is 7, 13, and 22 times greater than the corresponding changes in employment probabilities for workers 1, 2, or 3+ tracts away (within the same PUMA), 31, 53, and 78 times greater than for workers 1, 2, or 3+ PUMAs away (within the same state), and 273 and 3,068 times greater than for workers 1-2 and 3+ states away, respectively. In particular, the relative odds of changes in employment status for workers 3+ states away relative to workers in the local tract are 8 times higher than they were for the probability of obtaining a stimulus job.

The broader geographic incidence for general employment status is reflected in Figure 5 (Table A4, Col. 1), the analogue to Figure 3 (Table A2, Col. 1), which displays the share of the aggregate 500 job increase in employment attributable to workers initially employed in each distance bin relative to the target tract. Only 2.7% of the net employment change redounds to workers initially employed in the target tract, with 11% of the additional employment going to workers in other tracts within the PUMA, 29% to workers in other PUMAs within the target state, and a full 57% going to workers initially employed out of state.

Figure 6 (Table A3, Col. 2-9) illustrates how the change in employment probability for random workers in different distance bins varies across stimulus packages featuring new positions in different industry supersectors, while Figure 7 (Table A4, Col. 2-9) shows how the share of the net employment change enjoyed by areas defined by distance to the focal tract varies by the industry composition of the shock. Both figures show that the employment incidence across distance bins is very similar across industries, indicating that the geographic scope of labor markets is not wildly different across supersectors. Some small differences do exist: 15.7% and 45.2% of the change in net employment goes to workers within PUMA and within state, respectively, in stimuli featuring positions in the other services sector (which includes repair and maintenance, personal and laundry services, and

religious/civic organizations) sector, which is the sector featuring the most geographically concentrated employment incidence. By contrast, the corresponding figures are 12.6% and 41.1% for stimuli featuring positions in the retail/wholesale trade supersector (among the least geographically concentrated).

Figures 8 (Table A5) and 9 (Table A6) capture the corresponding heterogeneity in geographic incidence across stimulus simulations featuring positions at firms from different firm size quartile/firm average pay quartile combinations (averaging over industry supersectors). On average, stimulus packages featuring positions at large firms (4th quartile) with low average pay (2nd quartile) generate the most local incidence (2.9%, 15.4%, 46.0% percent of net employment gains within tract, within PUMA, and within state), and small firms (1st quartile) with high average pay (4th quartile) generate the least local incidence (2.3%, 12.7%, 40.3%).⁴⁰ Again, these differences are modest but nontrivial.

Our simulation procedure also generates counterfactual changes in worker mean utility necessary to clear the market for each origin job type following the various stimuli. Recall, though, that generating the market clearing allocation only requires computing standardized utility premia $\frac{C_g - C_{g'}}{\sigma}$. As discussed in section 3.2 and Appendix A3, we exploit the existence of a longer panel of years to generate estimates of σ that, given our assumption of a money-metric utility function, allow utility premia to be scaled in dollars of annual earnings. However, the assumptions that underlie the estimate of σ are stronger than for the relative joint surplus values (and are extremely unlikely to hold exactly). Thus, while proportional earnings changes for different skill and location categories ought to be reasonably well identified, the estimated dollar value of predicted welfare gains should be treated cautiously. Furthermore, since only relative utility changes are identified, we normalize the estimated utility impact on workers who are initially nonemployed 3 or more states away to be 0, so that all estimated utility changes are relative to this origin type.

Figure 10 (Table A7, Col. 1) provides the average utility impact (scaled in annual earnings equivalents) for random workers initially employed (or nonemployed) at different distances from the target tract for the “typical” stimulus package (again averaging over all 32 simulated stimuli featuring different firm compositions). The annual utility incidence is spread far more evenly across geographic space than the employment incidence. Workers initially working in the focal tract receive an estimated \$512 increase (in 2011 dollars) in money metric utility from the typical stimulus package (relative to workers 3+ states away), while workers initially working 1, 2, and 3 or more tracts away receive expected utility gains of \$211, \$164, and \$153 respectively. Workers initially working 1,2, and 3+ PUMAs away within the state receive the utility equivalent of \$140, \$115, and \$111 in annual earnings gains, while workers 1-2 states from the site of the shock receive gains of \$92 relative to

⁴⁰We chose the 2nd quartile of firm average pay to represent “low paying firms” rather than the 1st quartile so that our stimulus packages would be considered desirable for the receiving tract (most local development initiatives do not seek to increase the number of minimum wage jobs).

workers 3 or more states away. Figure 11 (Table A8, Col. 1) plots the share of total utility gains (relative to distant workers) that accrue to workers in each distance bin. Only 0.1% of total worker welfare gains accrue to workers within the focal tract, while over 81% accrue to workers initially employed out of state. Only about 2.5% of welfare gains accrue to workers within the PUMA associated with the focal tract. Thus, examining incidence from the perspective of welfare gains rather than employment gains suggests a far more geographically integrated labor market.

Figures 12 (Table A7, Col. 2-9) and 13 (Table A9, Col. 2-9) illustrate the heterogeneity in utility incidence by distance bin across stimuli featuring different industry supersectors and different firm size/firm average pay quartile combinations, respectively. Shocks featuring positions in the other services supersector generate the largest local welfare impact (\$632), while shocks featuring positions in the education and health supersector generate the smallest local impact (\$411). Stimuli featuring low paying positions at small firms generate the greatest local impact (\$546), while high paying positions at small firms generate the smallest local impact (\$459).

Another feature of the model is the ability to capture heterogeneity in shock incidence across workers in different skill classes, as proxied by initial employment status and earnings quintile. Figure 14 (Table A10, Col. 1) captures the share of the 500 job net employment gain that is enjoyed by workers whose initial earnings fall in each quintile in the national distribution, as well as workers who were nonemployed in the year prior to the simulated stimulus shock. 43.3% of the employment gains accrue to those initially nonemployed, while 15.3%, 14.3%, 11.3%, 8.6%, and 7.2% accrue to those at the 1st through 5th quintiles of the initial earnings distribution, respectively. The smaller values for initially high earning workers reflect the fact that such workers were less likely to transition to nonemployment in the absence of the shock. Figure 15 (Table A11) displays the share of the total worker welfare gain enjoyed by workers in each initial earnings quintile (and initially nonemployed workers). For the typical shock, 13.8% of utility gains accrue to initially nonemployed workers, while the share accruing to each earnings quintile increases in the level of initial earnings: 11.8%, 16.5%, 17.7%, 18.7%, and 21.1% for quintiles 1-5, respectively. These results suggest that existing high paid workers receive a disproportionate share of the welfare gains from a typical shock.

Figure 16 (Table A10, Col. 2-9) shows that there is relatively little heterogeneity across industry supersectors in the incidence of net employment gains among earnings quintiles; education/health and construction stimuli produce the greatest employment gains for the initially nonemployed, while retail/wholesale trade produces the least. Perhaps surprisingly, the firm size and particularly firm average pay quartiles of the firms generating the new positions (Figure 17, Table A12) are also predicted to have similar impacts on net employment gains across initial earnings quintile/nonemployment groups. Around 42% versus 44-45% of employment gains go to initially nonemployed workers from stimuli featuring

jobs at high paying versus low paying firms. By contrast, only $\sim 20\text{-}22\%$ of stimulus jobs at high paying firms are predicted to be obtained by initially non-employed workers versus $\sim 27\%$ percent for low paying firms, and $\sim 35\%$ versus 5% of jobs from high versus low paying firms go to workers initially in the highest earnings quintile, suggesting that the skill incidence of the actual stimulus jobs understates the employment gains that “trickle down” to initially nonemployed workers from labor demand shocks featuring a bias toward high skilled workers.

Figures 18 (Table A11, Col. 2-9) and 19 (Table A13) illustrate the importance of industry and firm size/firm pay heterogeneity for welfare incidence. There is very little heterogeneity in incidence across stimuli featuring different firm composition; even shocks featuring high paying jobs display only a slightly greater share of earnings gains ($\sim 1\%$) going to existing 5th earnings quintile workers.

We can also examine the degree to which the geographic scope of labor markets depends on the skill level. Figure 20 (Table A6) examines the change in employment probability for a randomly chosen worker whose initial job (or nonemployment) places him/her in a particular earnings quintile/distance bin combination. We see that nonemployed workers who most recently worked in the focal tract enjoy a large decrease in nonemployment rate of 6.2 percentage points, while the nonemployment decrease is only 0.5% and 0.3% for workers initially employed one or two census tracts away, indicating that the employment gains for existing nonemployed workers are particularly local. That said, employment gains decline with distance in a relatively similar fashion for all initial earnings quintiles. Existing nonemployed workers in the target district enjoy 2.1% of the total employment gains (Figure 12) despite only constituting 0.0005% of the population of potential workers.

Figure 21 (Table A14) displays the analogue to Figure 20 (Table A6) for welfare changes. The largest changes in utility from a typical shock, equivalent to an annual earnings gain of \$896, accrue to initially nonemployed workers located in the focal tract. Welfare changes are much smaller for local workers initially employed at the 1st quintile of earnings (\$363), and rise monotonically to \$545 for the 5th initial earnings quintile. Welfare gains decrease more quickly with distance from the focal tract for the higher income groups, however, creating rapid convergence in welfare gains across different income quintiles with distance from the focal tract. Thus, the results from Figure 15 (Table A11) that examined overall earnings incidence across skill levels that averaged across all distance categories obscured the much larger differences in earnings incidence between different skill levels that occurs among workers local to the shock.

Moreover, aggregating across all distance categories also obscured substantial heterogeneity across shocks featuring different firm composition in the welfare incidence among initial earnings categories for workers in the focal tract. Figure 22 (Table A15, Col 2-9) shows the money metric utility gains for only local workers by initial earnings quintile of the worker and industry supersector of the stimulus. Typical construction and other ser-

vices stimuli yield welfare gains for existing nonemployed workers equivalent to \$1,142 and \$1,196 in annual earnings, compared to \$702 and \$776 for stimuli featuring new jobs in the information or state/local government sectors. Workers in the highest (5th) initial earnings quintile reap expected utility gains of only \$387 from stimuli featuring jobs in the education/health supersector, while manufacturing and other services stimuli generate \$693 and \$649 for such workers (as noted above, stimuli in the other services supersector feature particularly local incidence for all skill levels).

Figure 23 (Table A16) shows the corresponding expected utility gains for workers in the focal tract by firm size/firm pay quartile combinations instead of industry. Stimuli featuring positions at small, high paying firms generate the least payoff for low skilled workers: \$670, \$282, and \$291 for nonemployed, 1st earnings quintile, and 2nd earnings quintile workers, respectively, while generating a substantial \$663 for 5th quintile workers. Large high paying firms (4th quartile of firm size) hire more locally for all skill levels, so that every earnings quintile initially employed (or most recently employed) in the focal tract earn substantially more from a shock featuring a large, high paying firm than a small, high paying firm. Interestingly, among low paying firms, larger size often generates smaller local gains. But stimuli featuring low paying firms (regardless of size) generate very large gains for nonemployed workers most recently employed in the focal tract (\$1,019 and \$1,051 for small and large low paying firms, respectively). The corresponding values for 1st and 2nd earnings quintiles are (\$433, \$410) and (\$466 and \$443) for small and large low paying firms. Thus, assuming similar impacts on rent price and feedback effects through the product market, it appears that the skill level of the positions being created matters a lot for incidence among skill classes for local workers, but much less for workers farther away.

Finally, the substantial heterogeneity in local skill incidence across industries and firm size/firm pay quartile combinations still misses further heterogeneity operating at the three-dimensional supersector/firm size/firm pay cell level. Figure 24 (Tables available upon request) plots the forecasted welfare gains by initial earnings level among workers initially in the focal tract for all 32 stimulus shock compositions that were simulated. There is a huge range of predicted gains. Welfare gains for initially nonemployed workers range from \$281 (small, high paying government positions) to \$1573 (small, low paying construction positions). For 1st earnings quintile workers, they range from \$184 (small, high paying information positions) to \$562 (small, low paying leisure/hospitality positions). For 5th quintile workers, they range from \$304 (large, low paying education/health positions) to small, high paying manufacturing positions). For local governments whose concern is primarily local incidence, these represent massive differences in the scale and skill intensity of the earnings incidence.

As noted in the estimation section, “plant relocation” shocks were also simulated that feature the same compositions of 500 new jobs as the stimulus packages but removed the jobs from a distant state (at least two states away from the focal tract) rather than from the

stock of “nonemployment” positions. However, since the “losing” locations are so far away from the “winning” tracts, and as shown above labor markets are still quite local and (to a lesser extent) regional, the employment and earnings incidence of such relocation shocks was virtually identical to their stimulus counterparts for locations within the winning state. Thus, I do not undertake a separate incidence analysis for our plant relocation shocks.

6.2 Natural Disasters

Recall that the “natural disaster” simulations remove at random 25%, 50%, or 100% of the destination jobs in the focal tract. Averaging over initial earnings categories, Figure 25 (Table A17) displays the increase in the probability of nonemployment for randomly chosen workers within different distance bins from the focal tract for each disaster intensity. Workers initially working (or nonemployed) in the focal tract experience increases in the probability of nonemployment in the destination year of 3.9%, 9.1%, and 24.5% from the 25%, 50%, and 100% disasters, respectively. The share of new nonemployment that falls upon workers initially in the local tract increases from 19.1% when 25% of local jobs disappear to 27.9% when 100% of local jobs disappear, suggesting that the employment incidence becomes increasingly geographically concentrated the more intense the local disaster (even when the disaster itself is in each case still contained within the same census tract). As with the stimulus shocks, in one sense these results suggest that labor markets remain very local: if mobility among labor markets were truly frictionless (and workers/positions were homogenous), so that the predicted employment incidence fell equally across all workers, then the expected share of lost jobs borne by local workers would be their share of the total workforce: 0.0039%. Thus, local workers experience a change in nonemployment probability that is over 8,000 times larger than it would be in a frictionless, homogenous world.

Note also that the employment incidence of the simulated disaster is more locally focused than for the stimulus packages, as measured by the local share of the total employment change. In the case of stimulus packages, most of the local workers would have been working (somewhere) in the absence of the shock (or are long-term nonemployed workers that would produce little joint surplus from employment, due to either preferences or low productivity), so that there was an effective limit to how local the employment incidence could be, thus forcing much of the net employment gain to be distributed across more distant locations. Analogously, since the nonemployment rate is well below 50% in most locations before the simulated shock, very few would have become nonemployed in the absence of the disaster, creating scope for a large effect. Indeed, the shocks produce a greater local labor surplus than existed nearly anywhere in the data; since most positions retain their existing workers (appearing in the model as a large increase in joint surplus from a match when the individual is an existing employee), it is very difficult for all of the local workers to find jobs. Thus, the model estimates reveal a natural asymmetry in the geographic scope of incidence between

positive and negative local demand shocks.

Figure 26 (Table A18) shows the average utility losses by distance bin for each disaster intensity. Expected utility losses (scaled as equivalent annual earnings losses) are severe for workers initially in the focal tract: \$-1,536, \$-2,802, and \$-4,653 for disasters featuring 25%, 50%, and 100% local job loss, respectively (relative to workers initially working 3+ states away). The welfare losses fall dramatically to \$-77, \$-147, and \$-278 for workers in an adjacent tract, and then decrease slowly in magnitude to \$-43, \$-83, and \$-151 for those initially working 1-2 states away from the focal tract. However, because within-tract workers are such a small share of the working population, the shares of aggregate worker welfare losses accounted for by the losses of workers initially employed in the focal tract are only 0.67%, 0.62%, and 0.57% for the three disaster intensities, respectively (Figure 27, Table A19) While this local share is 6 times larger than the local share of welfare gains for the stimulus packages, it is nonetheless trivial. As before, over 80% of the earnings incidence is predicted to fall on out-of-state workers. Again, we see that local shocks can have substantial impacts for local workers while still generating an overall incidence that is spread widely.

Figure 28 (Table A20) displays the share of all employment losses experienced by each employment status/initial earnings quintile. For disasters featuring 25% local job loss, 33% of lost net employment is experienced by those already nonemployed, with the share falling monotonically from 17.3% to 8.5% as one moves from the 1st to the 5th initial earnings quintile. Thus, high skilled workers seem relatively well insulated from employment losses, instead taking jobs from those at lower skill levels, creating a cascade of sorts. However, as the disaster becomes more intense, the burden of employment loss becomes more equally shared, with only 28% accounted for by initially nonemployed, and 11% accounted for by those initially in the highest earnings quintile.

However, Figure 29 (Table A21), which examines employment incidence by distance and initial earnings jointly, paints a richer picture. Among those initially employed in the focal tract, the increase in the probability of nonemployment from the least severe (25%) disaster is actually larger for employed workers than for initially nonemployed workers: initially nonemployed workers experience a 1.8 percentage point increase in destination nonemployment, while workers at initial earnings quintiles 1-5 experience increases in nonemployment rate of 5.9, 4.8, 4.2, 3.8, and 2.9 percentage points, respectively. This is primarily due to the fact that initially nonemployed workers had the least to lose: they were fairly likely to be nonemployed again in the absence of a disaster. However, among workers a tract away or further, the employment losses are greatest among the existing nonemployed. As the disaster becomes more severe, this pattern becomes even more pronounced. For the most severe (100% job loss) disaster, initially nonemployed local workers experience a 5.2 percentage point increase in nonemployment rate, while local 1st and 5th earnings quintile workers experience 30.4% and 23.9 percentage point increases, respectively (Figure 30 (Table A22)).

Figure 31, Table A23) displays the share of all worker welfare losses experienced by each employment status/initial earnings quintile. The shares are almost identical across disaster intensities, with 14% falling on nonemployed workers, and 12%, 17%, 18%, 19%, and 21% falling on workers in earnings quintiles 1-5, respectively. As with employment incidence, however, these numbers obscure substantial variation in the relative skill incidence of disasters by distance from the focal site. For disasters involving a 25% job loss, workers 1-2 states away experience utility losses, relative to workers 3+ states away, equivalent to between \$-41 and \$-43 in annual earnings regardless of initial skill incidence (Figure 32, Table A24). The values are between \$-145 and \$-152 for the 100% job loss disasters (Figure 33, Table A25). Differences in welfare losses are similarly small for all distance bins except workers initially employed (or nonemployed) in the focal tract. However, workers initially nonemployed within the focal tract are predicted to lose the equivalent of \$146 in utility in the 25% disaster, while workers in earnings quintiles 1-5 are predicted to lose \$-1,249, \$-1,552, \$-1,746, \$-1,961, and \$-2,015, respectively. For the disasters featuring 100% local job loss, these values rise to \$-431 and \$-3,618, \$-4,526, \$-5,148, \$-6,150, and \$-6,291 respectively. Thus, welfare losses are particularly large among high skilled workers (who had the most to lose), although smaller as a share of initial utility from annual earnings.

Finally, while quantifying the employment and utility incidence of disasters is important for allocating relief funds, policymakers and local communities are also worried about being inundated by flows of migrants away from disaster sites. Thus, Figure 34 (Table A26) displays, for each disaster intensity, the change in the probability of being employed at firms in each distance bin relative to the focal tract for workers who were initially employed (or nonemployed) in the census tract hit by the natural disaster. First, note that for the mildest disaster, the decrease in within-tract employment for workers initially at the focal tract is only 10%, despite a 25% overall decrease in local positions. This is in part because many of these workers would have moved to jobs away from the tract even in the absence of the shock, but also because existing local workers are able to retain a greater share of the jobs that remain. Even when all local jobs are lost, the decrease in the share of local workers staying in the tract is only 66.6%, revealing that 1/3 of such workers would have moved out in absence of the disaster. Adjacent tracts absorb an extra 0.4%, 1.0%, and 2.5% of workers initially in the focal tract, respectively, in the 25%, 50%, and 100% job loss scenarios, relative to a counterfactual in which no disaster occurs. Overall, an additional 1.8%, 4.3%, and 11.0% of the workers initially in the focal tract end up employed in other tracts within the original PUMA after the 25%, 50%, and 100% job loss scenarios. Locations outside the PUMA but within the state take on an additional 2.3%, 5.6%, and 15.4% of those initially employed in the focal tract in the three disaster scenarios, while locations outside the state take on an additional 2.3%, 5.6%, and 15.7% in the three scenarios, with the remaining share of workers experiencing nonemployment. Thus, while a relatively small share of employees find employment nearby, this share increases in the degree of initial

displacement.

Figure 35 (Table A27) displays separate distributions of destination employment locations for target tract workers in each employment status/initial earnings quintile. Even in the most severe disaster, only an additional 1.8% of initially nonemployed workers move away from the focal tract, relative to the counterfactual. The few that would have gotten local jobs remain unemployed instead. By contrast, the share moving to nearby locations is much larger for initially employed workers, and is increasing in the skill level of the worker. Since most high paid workers are retained or continue to work nearby in the absence of the shock (only 16% would have transitioned away from the focal tract), the extreme 100% disaster engenders a particularly large mobility response for such workers: an additional 14.2% move to another tract within the PUMA (relative to the counterfactual), an additional 20.6% move to a different PUMA within the state, and an additional 23.9% move to a different state. For workers initially in the 1st earnings quintile, who were more mobile in the absence of the disaster, the corresponding increases are only 9.2%, 13.1%, and 10.0%, respectively. Relative to lower skilled workers, the mobility response for initially high earning workers is disproportionately muted for lesser disasters, though, because they are better able to capturing the remaining local jobs than less skilled workers (Figure 36, Table A28).

7 Conclusion

Building on the approach of Choo and Siow (2006), this paper models the transition of the U.S. labor market across adjacent years as a large-scale assignment game with transferable utility, and uses a very large set of estimated parameters from the model to simulate the welfare incidence across locations and worker skill categories of a variety of alternative local labor demand shocks designed to resemble different stimulus packages and natural disasters.

I show that allowing unobserved heterogeneity at the fundamental worker-position match level (rather than each side only holding preferences over types as in Choo and Siow (2006) and Galichon and Salanié (2015)) can still be used to produce forecasts of welfare incidence on both sides of the market from changes in agent type composition on either side of the market, even when singles are either not observed or observed on only one side of the market. By basing simulations on millions of composite joint surplus parameters rather than reducing the data to a much smaller set of fundamental utility or production function parameters, the “sufficient statistics” approach used here can fully exploit the massive scale of the administrative LEHD database to capture multidimensional heterogeneity on both sides of a two-sided market without placing unjustified structure on the job matching technology.

The method can be customized to forecast the incidence of any particular shock composition or magnitude in any location, and incidence can be determined across groups of

agents on either side of the market defined by any arbitrary combination of observed characteristics, including categorical characteristics without a natural ordering such as race, industry or location. Given appropriate administrative matching data, the approach here could also be easily adapted to the student-college matching or patient-doctor matching contexts, among other applications.

I find that U.S. labor markets are still quite local, in that the per-worker welfare gains from a locally targeted labor demand shock are substantially larger for workers in the focal census tract than even workers one or two tracts away. Nonetheless, because the workers initially working within a very small radius of the local shock are such a small share of the entire U.S. labor force competing for positions, I also find that in most specifications greater than 80% of the welfare gain from a very local stimulus package, regardless of firm composition, redounds to workers initially working out of state, with only about 0.1% of the welfare gains going to existing workers in the focal census tract.

I also document a high degree of heterogeneity in skill incidence among very local workers across demand shocks featuring different firm size, firm average pay, and industry super-sector composition, suggesting that the type of firm targeted by a local development policy has major implications for the groups of workers most likely to benefit. That said, as these alternative shocks ripple across space through a chain of job transitions, their skill incidence becomes increasingly similar, so that the overall skill composition of welfare gains across all workers (not just local workers) is extremely similar across different types of demand shocks. Thus, local policymakers need to be much more careful about which projects are subsidized than would state or national funders of local projects, since such funders would internalize the ripple effects.

Finally, I find that positive and negative shocks have asymmetric impacts, with negative shocks displaying a much greater geographic concentration of employment losses than the corresponding employment gains from positive shocks. This is because most workers would have been working anyway in the absence of a positive shock, while most have jobs at risk when negative shocks occur.

Going forward, two extensions seem particularly worthwhile. First, following Caliendo et al. (2015), rather than computing incidence over a one year horizon, the assignment game could be played several times in a row, updating the initial assignments after each simulated year-to-year transition, so that the long-run welfare incidence of labor demand and supply shocks could be evaluated (though this might require explicitly modeling expectations of shock duration). Second, following Galichon and Salanié (2015) and Chiappori et al. (2009), one could relax the assumption that the idiosyncratic job-match-level surplus shocks are i.i.d. across worker and position types so as to incorporate an even more flexible set of substitution patterns than those featured in the model estimated here.

References

- Abowd, John M, Bryce E Stephens, Lars Villhuber, Fredrik Andersson, Kevin L McKinney, Marc Roemer, and Simon Woodcock (2009) ‘The lehd infrastructure files and the creation of the quarterly workforce indicators.’ In ‘Producer Dynamics: New Evidence from Micro Data’ (University of Chicago Press) pp. 149–230
- Autor, David H, David Dorn, Gordon H Hanson, and Jae Song (2014) ‘Trade adjustment: Worker-level evidence.’ *The Quarterly Journal of Economics* 129(4), 1799–1860
- Busso, Matias, Jesse Gregory, and Patrick Kline (2013) ‘Assessing the incidence and efficiency of a prominent place based policy.’ *The American Economic Review* 103(2), 897–947
- Caliendo, Lorenzo, Maximiliano Dvorkin, and Fernando Parro (2015) ‘The impact of trade on labor market dynamics.’ Technical Report, National Bureau of Economic Research
- Chiappori, Pierre-Andr, and Bernard Salani (2016) ‘The econometrics of matching models.’ *Journal of Economic Literature* 54(3), 832–61
- Chiappori, Pierre-Andre, Murat Iyigun, and Yoram Weiss (2009) ‘Investment in schooling and the marriage market.’ *The American Economic Review* 99(5), 1689–1713
- Choo, Eugene (2015) ‘Dynamic marriage matching: An empirical framework.’ *Econometrica* 83(4), 1373–1423
- Choo, Eugene, and Aloysius Siow (2006) ‘Who marries whom and why.’ *Journal of political Economy* 114(1), 175–201
- Decker, Colin, Elliott H Lieb, Robert J McCann, and Benjamin K Stephens (2013) ‘Unique equilibria and substitution effects in a stochastic model of the marriage market.’ *Journal of Economic Theory* 148(2), 778–792
- Freedman, Matthew (2013) ‘Targeted business incentives and local labor markets.’ *Journal of Human resources* 48(2), 311–344
- Galichon, Alfred, and Bernard Salanié (2015) ‘Cupids invisible hand: Social surplus and identification in matching models’
- Galichon, Alfred, Bernard Salanié et al. (2017) ‘The econometrics and some properties of separable matching models.’ *American Economic Review* 107(5), 251–255
- Glaeser, Edward L, Joshua D Gottlieb et al. (2008) ‘The economics of place-making policies.’ *Brookings Papers on Economic Activity* 39(1 (Spring)), 155–253
- Greenstone, Michael, Richard Hornbeck, and Enrico Moretti (2010) ‘Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings.’ *Journal of Political Economy* 118(3), 536–598
- Gregory, Jesse (2013) ‘The impact of post-katrina rebuilding grants on the resettlement choices of new orleans homeowners.’ *Work. Pap., Univ. Wisconsin–Madison*
- Kennan, John, and James R Walker (2011) ‘The effect of expected income on individual migration decisions.’ *Econometrica* 79(1), 211–251

- Kline, Patrick, and Enrico Moretti (2013) ‘Place based policies with unemployment.’ Technical Report, National Bureau of Economic Research
- Koopmans, Tjalling C, and Martin Beckmann (1957) ‘Assignment problems and the location of economic activities.’ *Econometrica: journal of the Econometric Society* pp. 53–76
- LeGower, Michael, and Randall Walsh (2017) ‘Promise scholarship programs as place-making policy: Evidence from school enrollment and housing prices.’ *Journal of Urban Economics*
- Manning, Alan, and Barbara Petrongolo (2017) ‘How local are labor markets? evidence from a spatial job search model.’ *American Economic Review*
- McCall, John Joseph (1970) ‘Economics of information and job search.’ *The Quarterly Journal of Economics* pp. 113–126
- Menzel, Konrad (2015) ‘Large matching markets as two-sided demand systems.’ *Econometrica* 83(3), 897–941
- Monte, Ferdinando, Stephen J Redding, and Esteban Rossi-Hansberg (2015) ‘Commuting, migration and local employment elasticities.’ Technical Report, National Bureau of Economic Research
- Moretti, Enrico (2010) ‘Local multipliers.’ *The American Economic Review* pp. 373–377
- Neumark, David, and Helen Simpson (2014) ‘Place-based policies.’ Technical Report, National Bureau of Economic Research
- Nimczik, Jan Sebastian (2017) ‘Job mobility networks and endogenous labor markets’
- Roback, Jennifer (1982) ‘Wages, rents, and the quality of life.’ *Journal of political Economy* 90(6), 1257–1278
- Roth, Alvin E, and Marilda Sotomayor (1992) ‘Two-sided matching.’ *Handbook of game theory with economic applications* 1, 485–541
- Sattinger, Michael (1993) ‘Assignment models of the distribution of earnings.’ *Journal of economic literature* 31(2), 831–880
- Schmutz, Benoît, and Modibo Sidibe (2016) ‘Frictional labor mobility’
- Shapley, LS, and Martin Shubik (1972) *Game Theory in Economics: Chapter 3: the” rules of the Games”* (Rand)

Table 1: Specifications for Alternative Labor Demand Shocks Used in Counterfactual Simulations

Spec. No.	Number of Jobs (or % of Tract's Jobs)	Firm Avg. Earn. Quartile	Firm Size Quartile	Industry Supersector	Shock Type
1	500	2	1	Information	Stimulus
2	500	2	4	Information	Stimulus
3	500	4	1	Information	Stimulus
4	500	4	4	Information	Stimulus
5	500	2	1	Manufacturing	Stimulus
6	500	2	4	Manufacturing	Stimulus
7	500	4	1	Manufacturing	Stimulus
8	500	4	4	Manufacturing	Stimulus
9	500	2	1	Trade/Trans./Utilities	Stimulus
10	500	2	4	Trade/Trans./Utilities	Stimulus
11	500	4	1	Trade/Trans./Utilities	Stimulus
12	500	4	4	Trade/Trans./Utilities	Stimulus
13	500	2	1	Other Services	Stimulus
14	500	2	4	Other Services	Stimulus
15	500	4	1	Other Services	Stimulus
16	500	4	4	Other Services	Stimulus
17	500	2	1	Education & Health	Stimulus
18	500	2	4	Education & Health	Stimulus
19	500	4	1	Education & Health	Stimulus
20	500	4	4	Education & Health	Stimulus
21	500	2	1	Leisure & Hospitality	Stimulus
22	500	2	4	Leisure & Hospitality	Stimulus
23	500	4	1	Leisure & Hospitality	Stimulus
24	500	4	4	Leisure & Hospitality	Stimulus
25	500	2	1	Government	Stimulus
26	500	2	4	Government	Stimulus
27	500	4	1	Government	Stimulus
28	500	4	4	Government	Stimulus
29	500	2	1	Construction	Stimulus
30	500	2	4	Construction	Stimulus
31	500	4	1	Construction	Stimulus
32	500	4	4	Construction	Stimulus
33	500	2	1	Information	Relocation
34	500	2	4	Information	Relocation
35	500	4	1	Information	Relocation
36	500	4	4	Information	Relocation
37	500	2	1	Manufacturing	Relocation
38	500	2	4	Manufacturing	Relocation
39	500	4	1	Manufacturing	Relocation
40	500	4	4	Manufacturing	Relocation
41	500	2	1	Trade/Trans./Utilities	Relocation
42	500	2	4	Trade/Trans./Utilities	Relocation
43	500	4	1	Trade/Trans./Utilities	Relocation
44	500	4	4	Trade/Trans./Utilities	Relocation
45	25%	All	All	All	Relocation
46	50%	All	All	All	Relocation
47	100%	All	All	All	Relocation

Figure 1: Probability of Obtaining Stimulus Job by Distance From Focal Tract: Average across All Simulated Stimuli

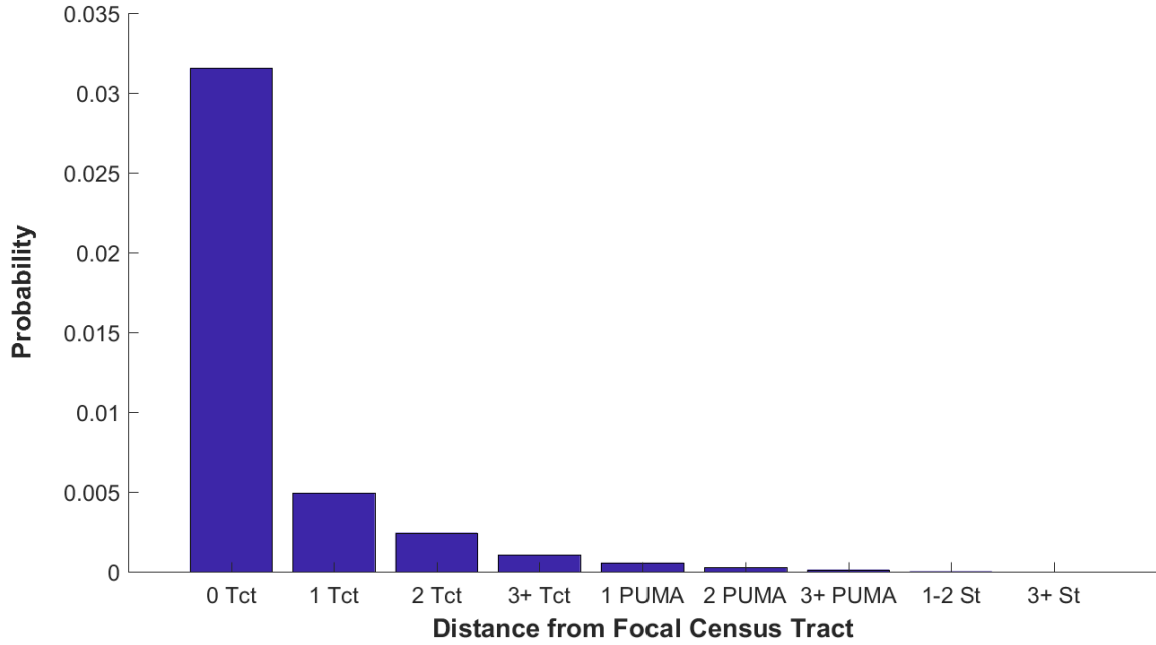


Figure 2: Fraction of Workers in Each Distance Bin




Figure 3: Share of Stimulus Jobs Obtained by Workers at Each Distance From Focal Tract: Average across All Simulated Stimuli

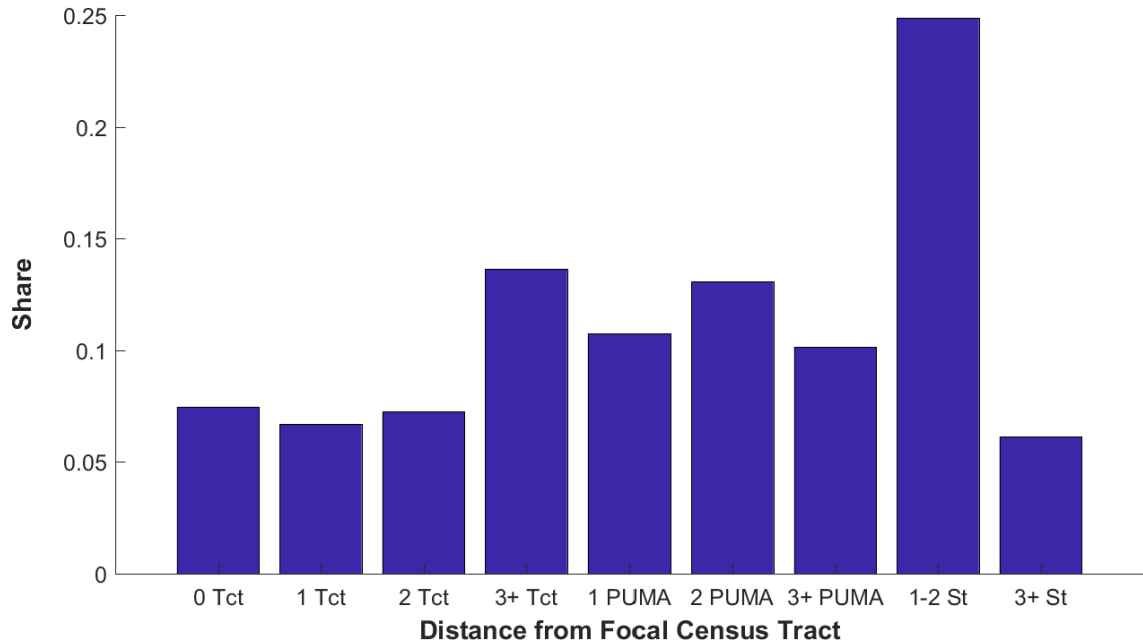


Figure 4: Change in P(Employed) by Distance From Focal Tract: Average across All Simulated Stimuli




Figure 5: Share of Additional Employment Obtained by Workers at Each Distance From Focal Tract: Average across All Simulated Stimuli

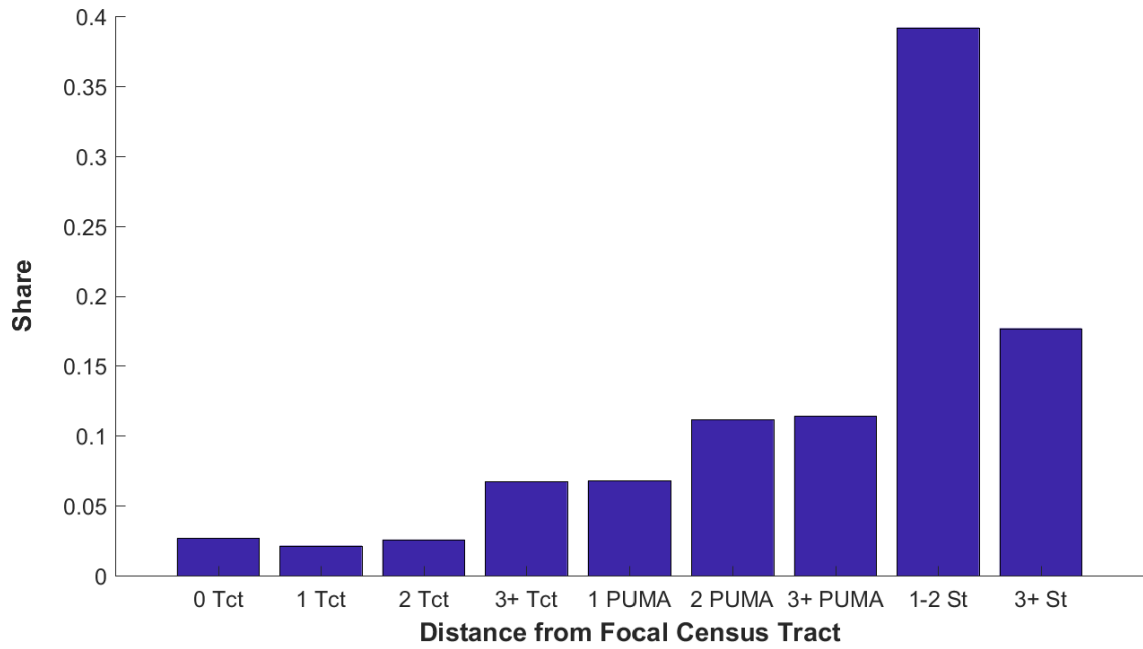


Figure 6: Change in P(Employed) by Distance From Focal Tract and Industry of Stimulus

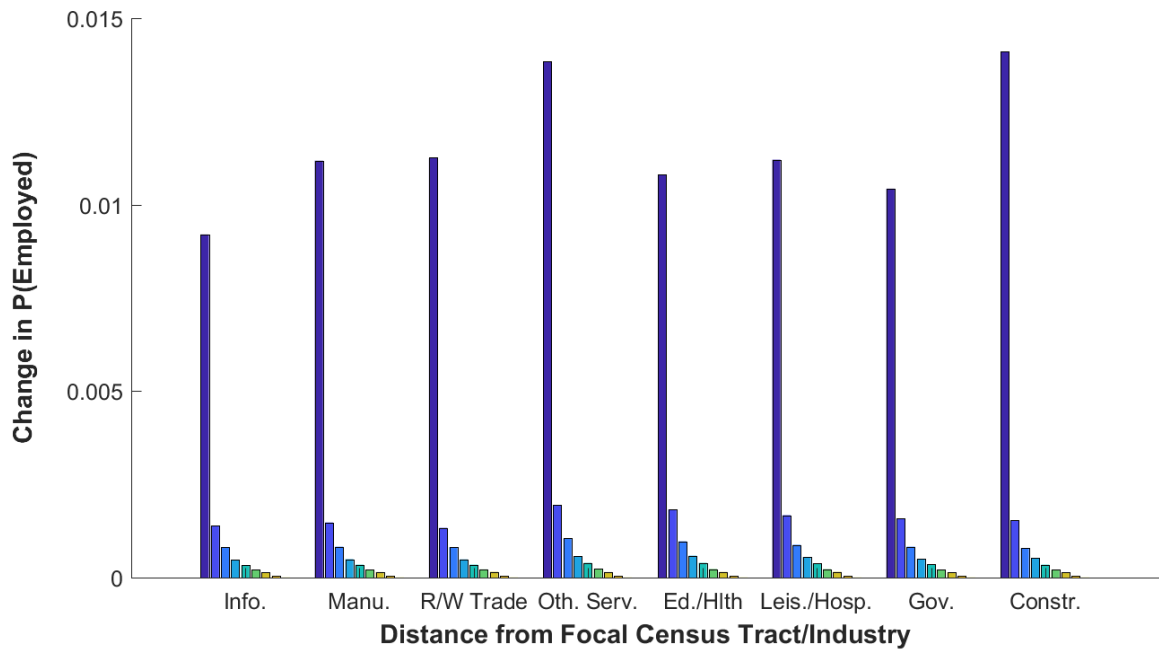


Figure 7: Share of Additional Employment Obtained by Workers by Distance From Focal Tract and Industry Composition of the Stimulus Package

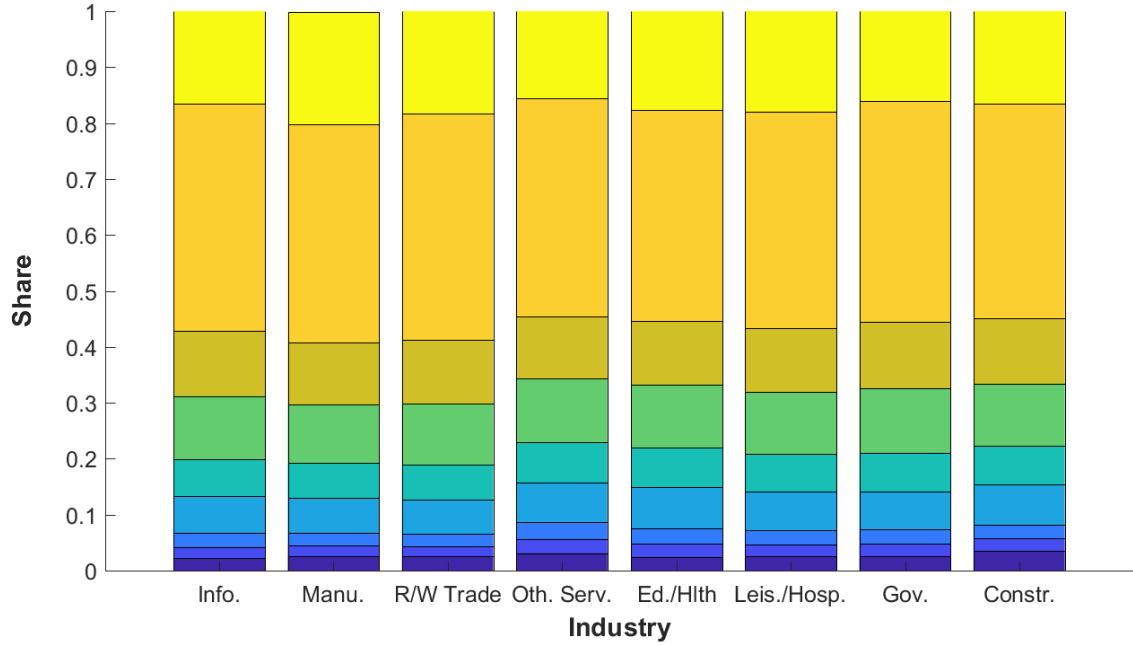


Figure 8: Change in P(Employed) by Distance From Focal Tract and Firm Size/Firm Average Pay Composition of the Stimulus Package

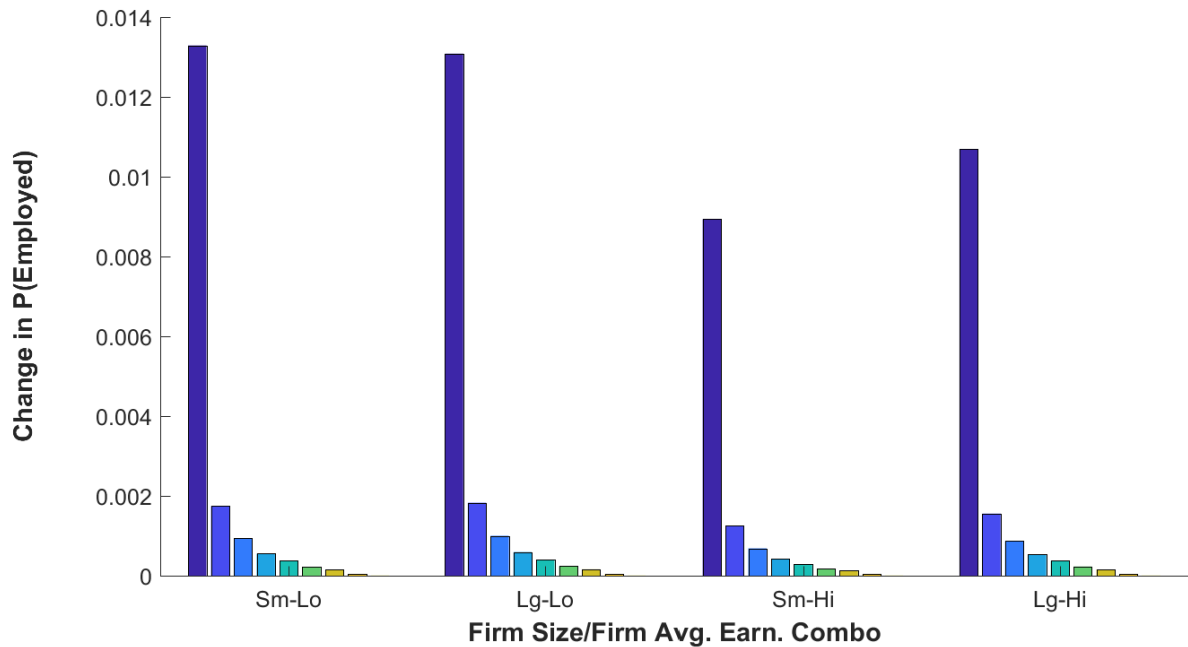


Figure 9: Share of Additional Employment Obtained by Workers by Distance From Focal Tract and Firm Size/Firm Average Pay Composition of the Stimulus Package

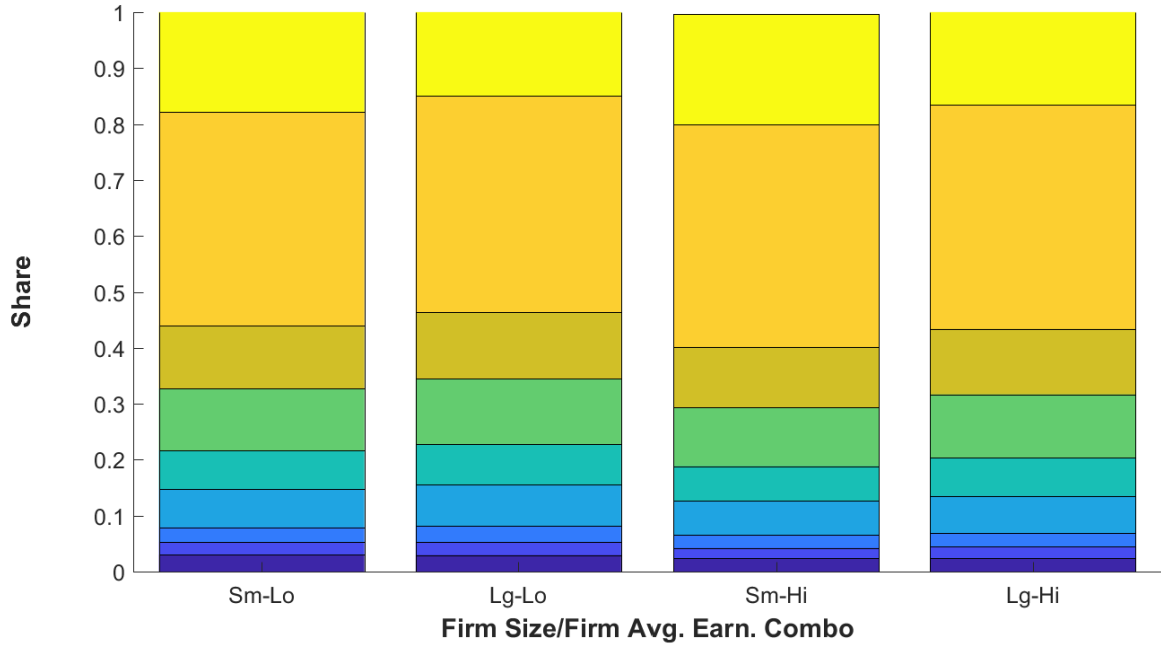


Figure 10: Expected Utility Changes by Distance From Focal Tract: Average across All Simulated Stimuli

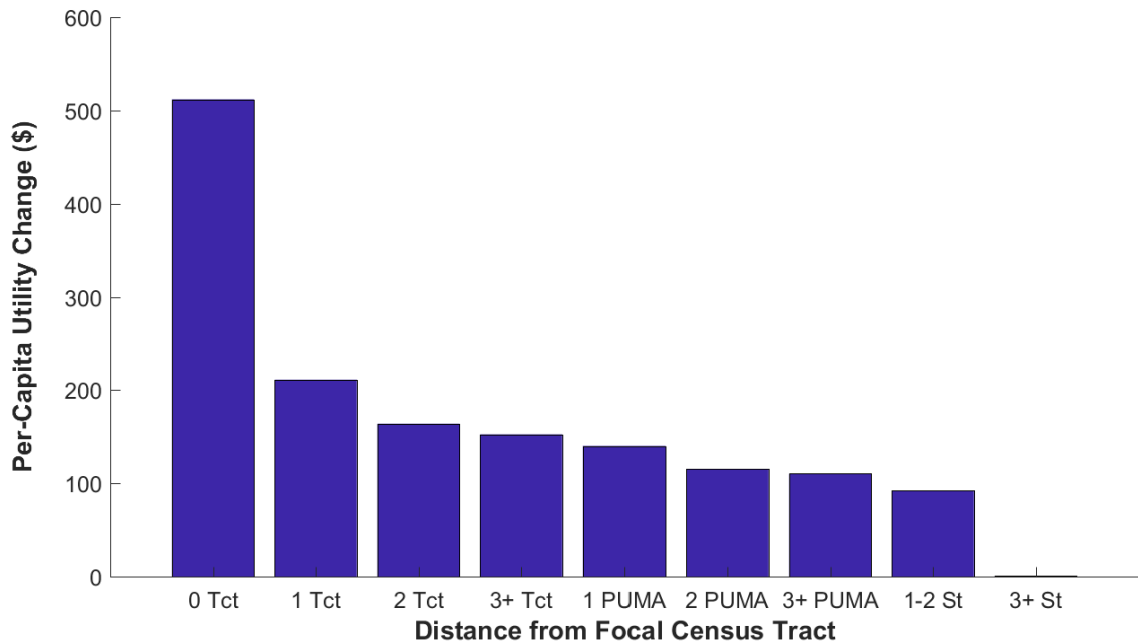


Figure 11: Share of Total Utility Gains by Distance From Focal Tract: Average across All Simulated Stimuli

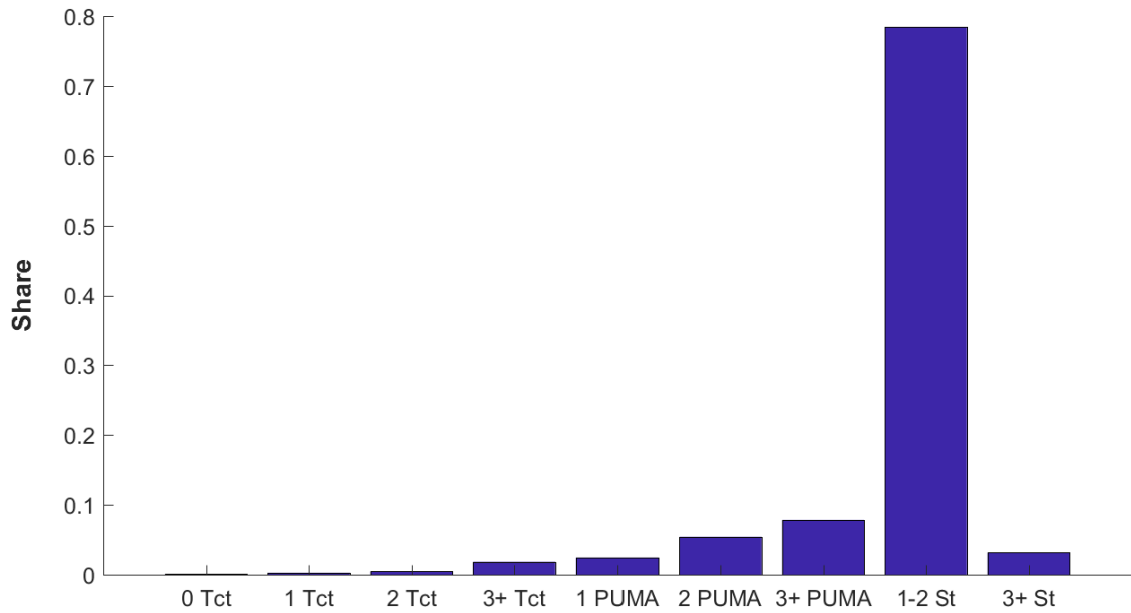


Figure 12: Expected Utility Changes by Distance From Focal Tract and Industry Composition of the Stimulus Package

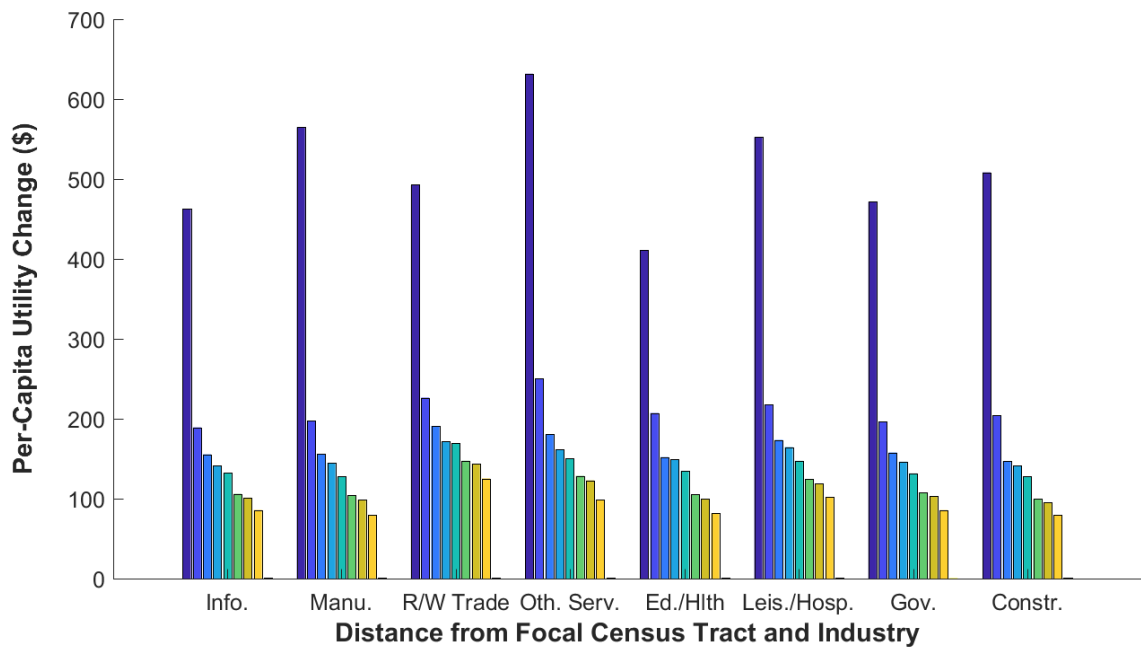


Figure 13: Expected Utility Changes by Distance From Focal Tract and Firm Size/Firm Average Pay Composition of the Stimulus Package

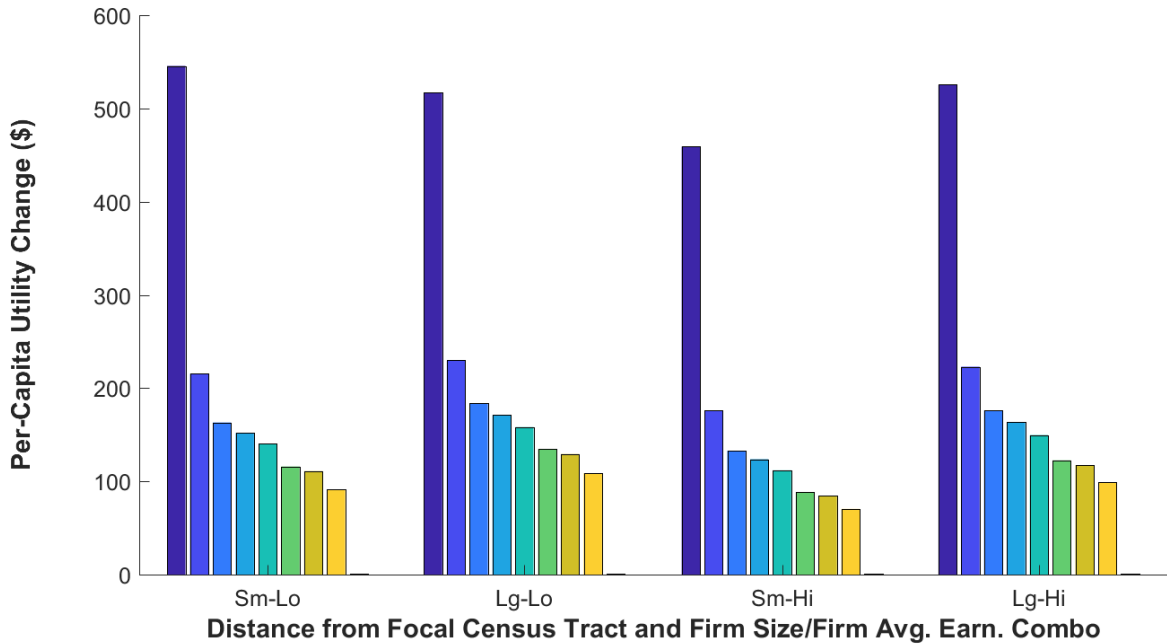


Figure 14: Share of Additional Employment among Workers Initially Employed (or Nonemployed) at Different Initial Earnings Quintiles (or Nonemployment): Average across All Simulated Stimuli

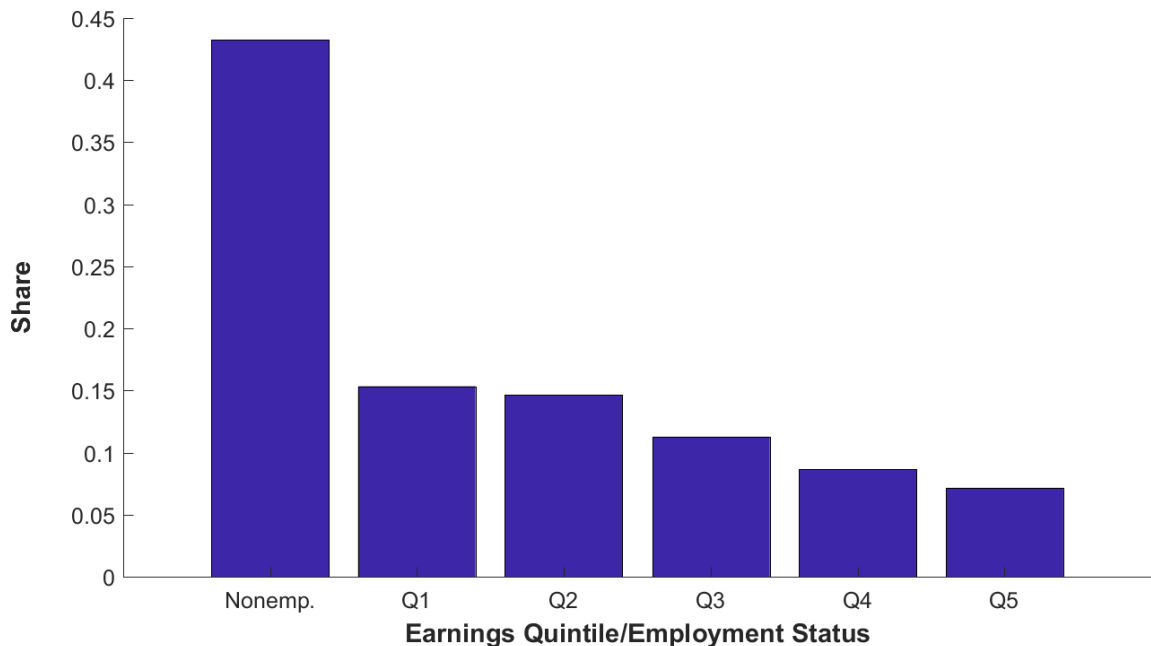


Figure 15: Share of Total Utility Gains among Workers Initially Employed (or Nonemployed) at Different Initial Earnings Quintiles (or Nonemployment): Average across All Simulated Stimuli

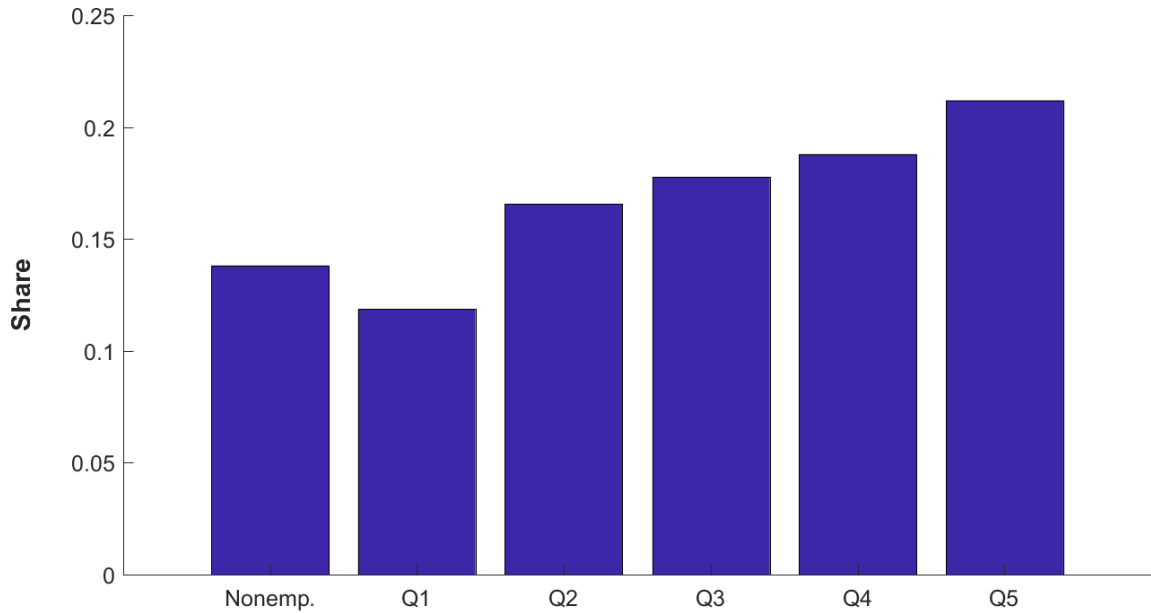


Figure 16: Share of Additional Employment among Workers Initially Employed (or Nonemployed) at Different Initial Earnings Quintiles (or Nonemployment) by Industry Composition of the Stimulus Package

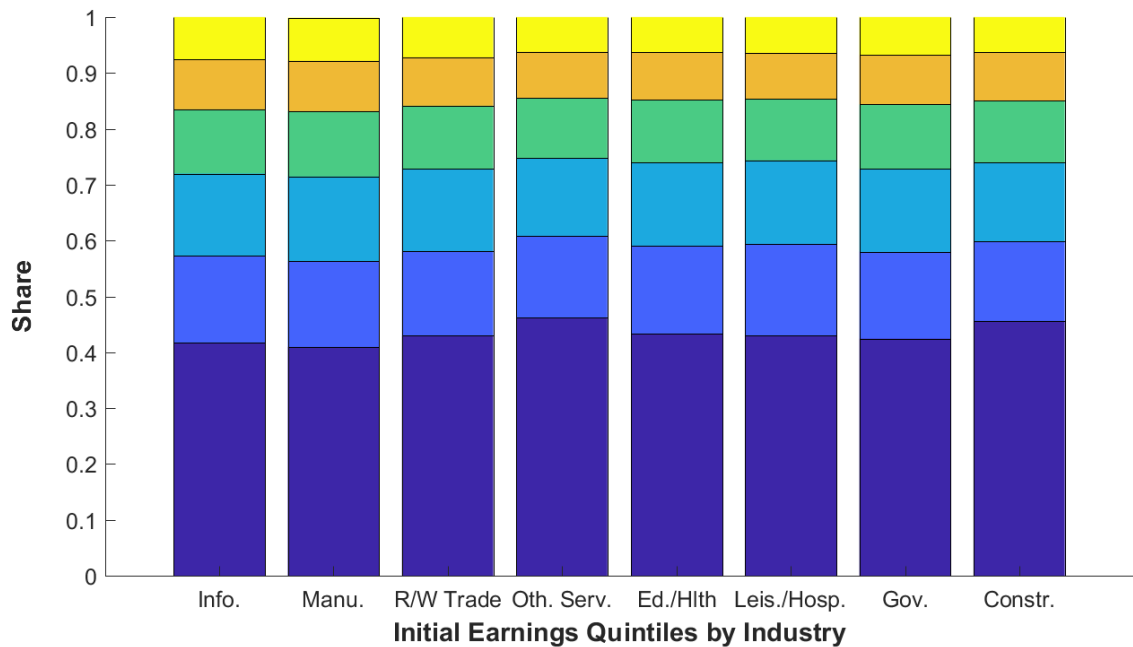


Figure 17: Share of Additional Employment among Workers Initially Employed (or Nonemployed) at Different Initial Earnings Quintiles (or Nonemployment) by Firm Size/Firm Average Pay Composition of the Stimulus Package

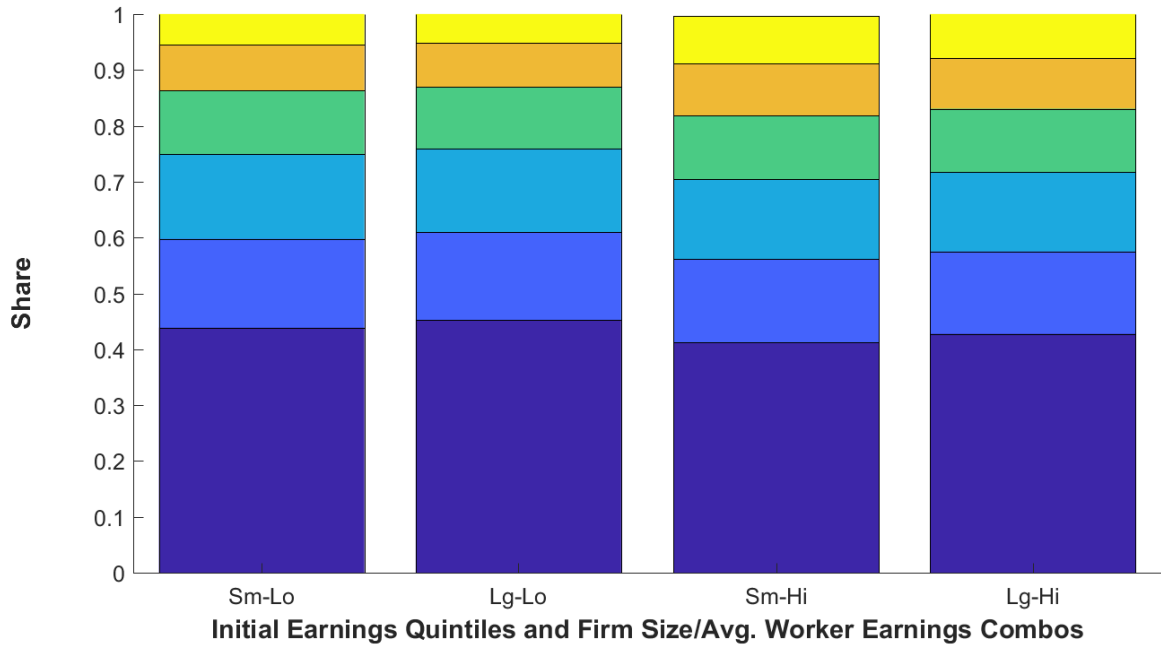


Figure 18: Share of Total Utility Gains among Workers Initially Employed (or Nonemployed) at Different Initial Earnings Quintiles (or Nonemployment) by Industry Composition of the Stimulus Package




Figure 19: Share of Total Utility Gains among Workers Initially Employed (or Nonemployed) at Different Initial Earnings Quintiles (or Nonemployment) by Firm Size/Firm Average Pay Composition of the Stimulus Package




Figure 20: Change in P(Employed) Among Workers Initially Employed at Different Combinations of Initial Earnings Quintile (or Nonemployed) and Distance from Focal Tract: Average across All Simulated Stimuli

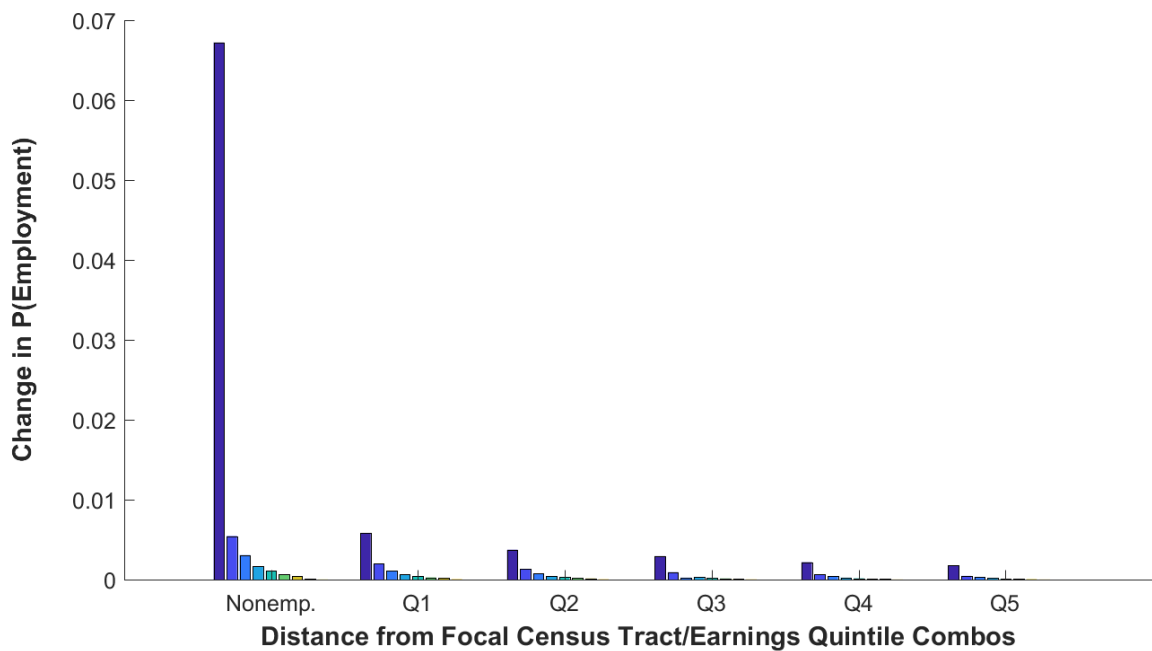


Figure 21: Expected Utility Changes Among Workers Initially Employed at Different Combinations of Initial Earnings Quintile (or Nonemployed) and Distance from Focal Tract: Average across All Simulated Stimuli

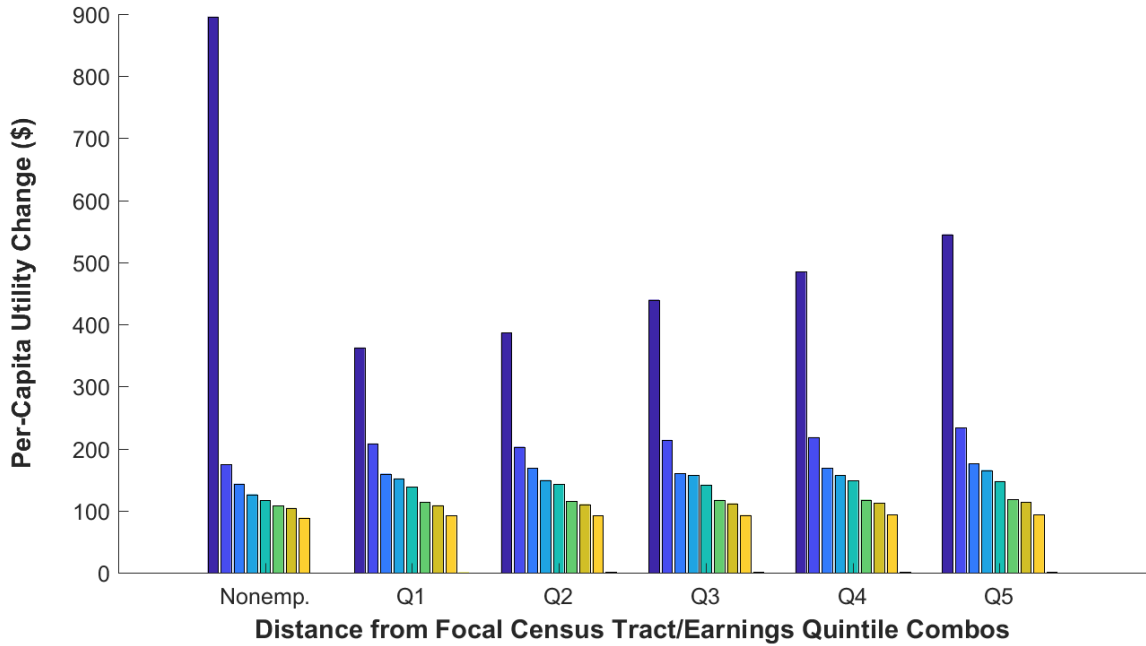


Figure 22: Expected Utility Changes Among Workers Originally Working in the Targeted Tract by Initial Earnings Quintile: By Industry Supersector

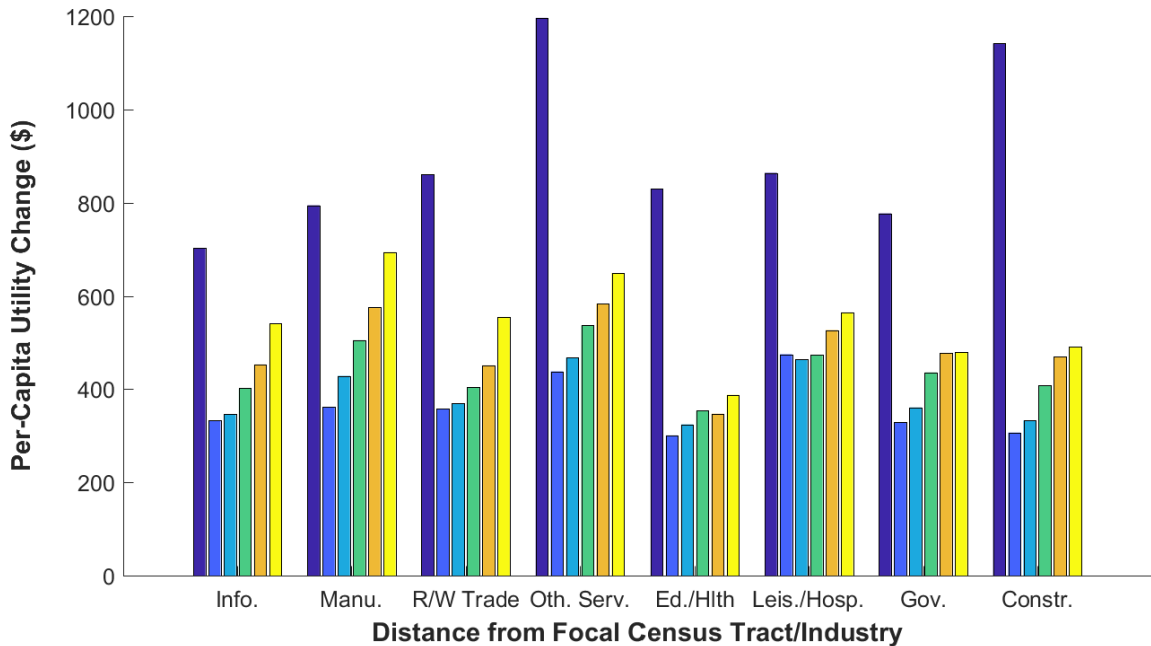


Figure 23: Expected Utility Changes Among Workers Originally Working in the Targeted Tract by Initial Earnings Quintile: By Firm Size Quartile/Firm Pay Quartile Combination

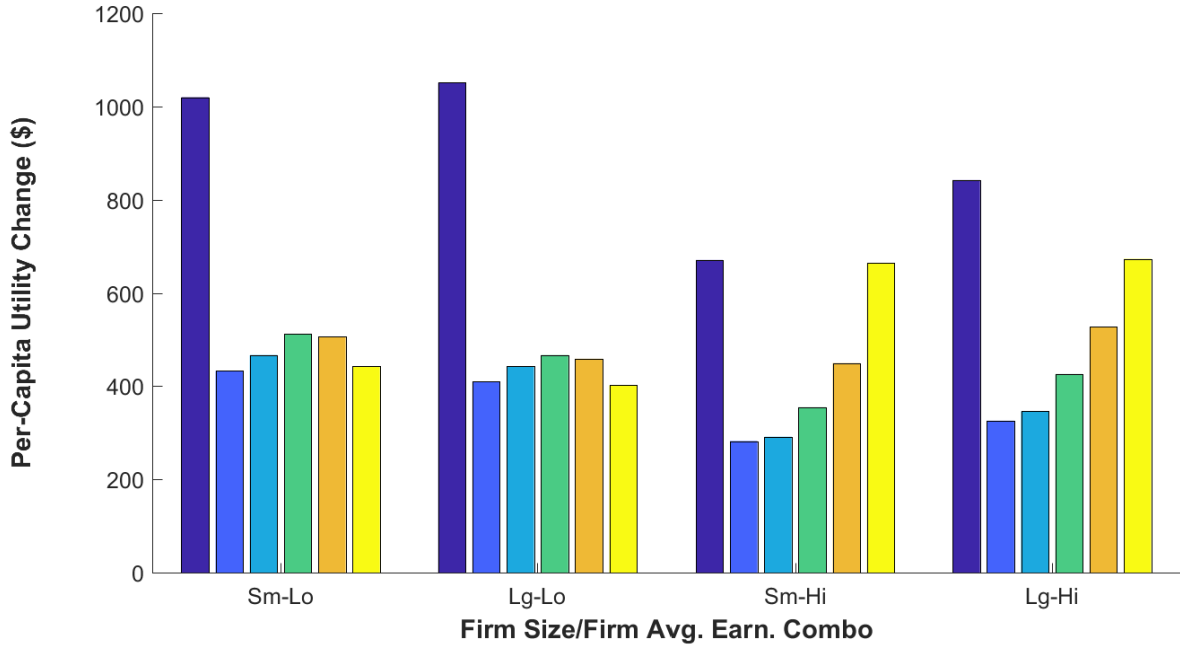


Figure 24: Expected Utility Changes Among Workers Originally Working in the Targeted Tract by Initial Earnings Quintile: All Stimulus Packages

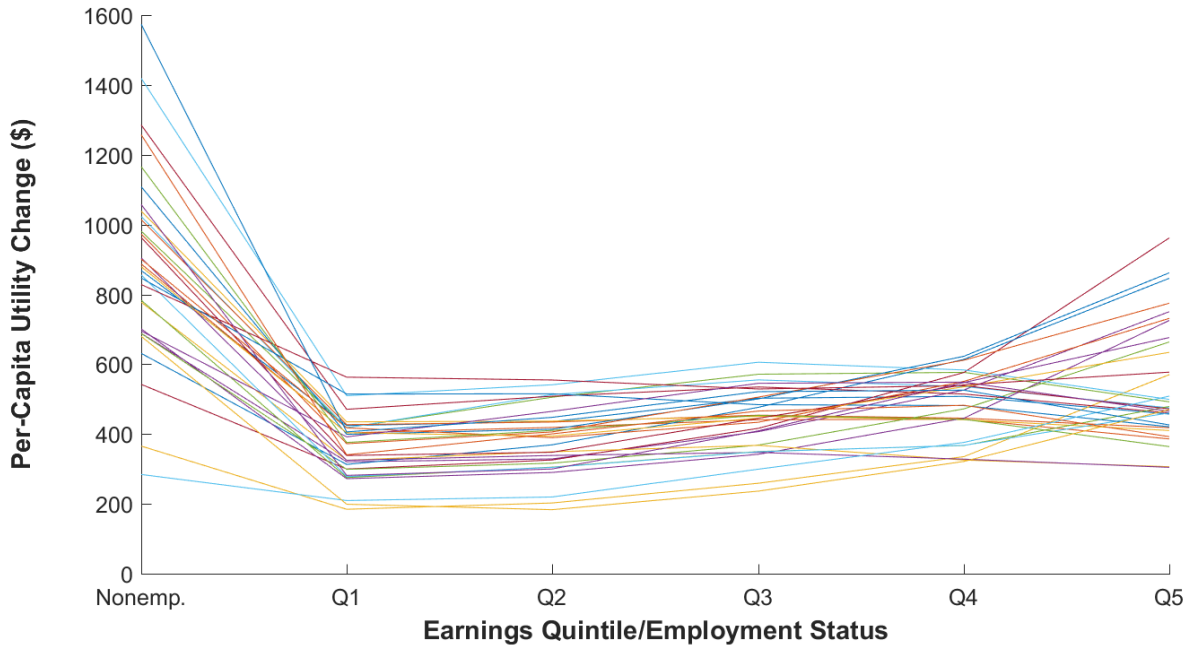


Figure 25: Change in $P(\text{Unemployed})$ from a Natural Disaster by Distance From Focal Tract and Severity of the Disaster (25%/50%/100% Jobs Lost)

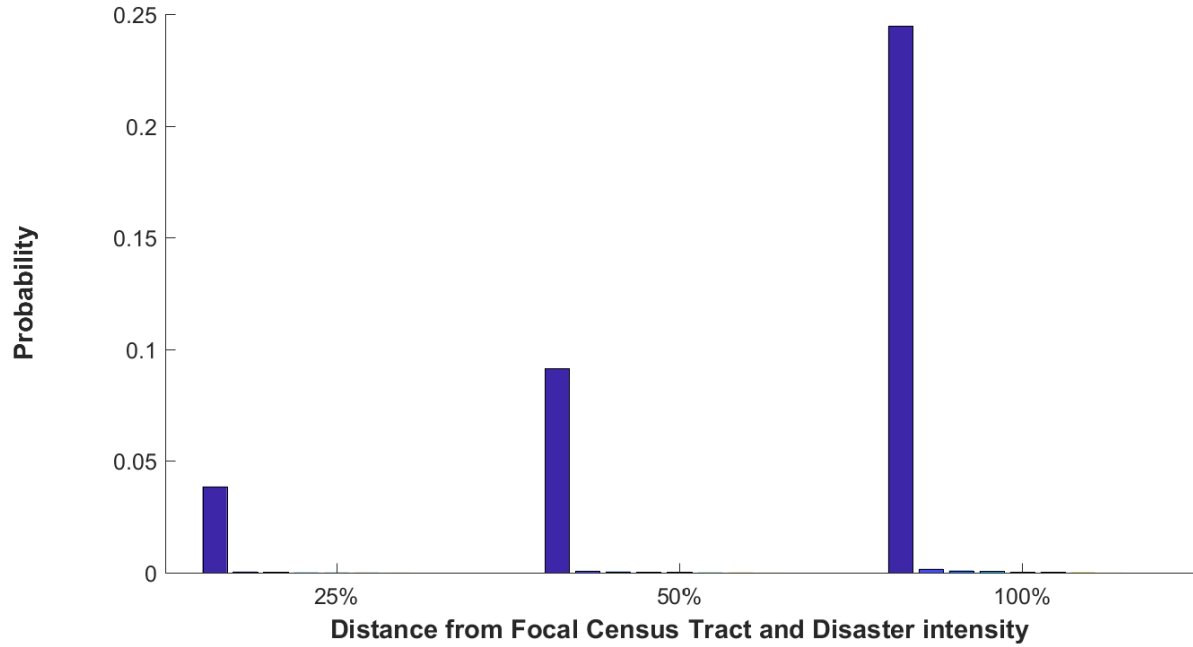


Figure 26: Expected Utility Loss from a Natural Disaster by Distance From Focal Tract and Severity of the Disaster (25%/50%/100% Jobs Lost)

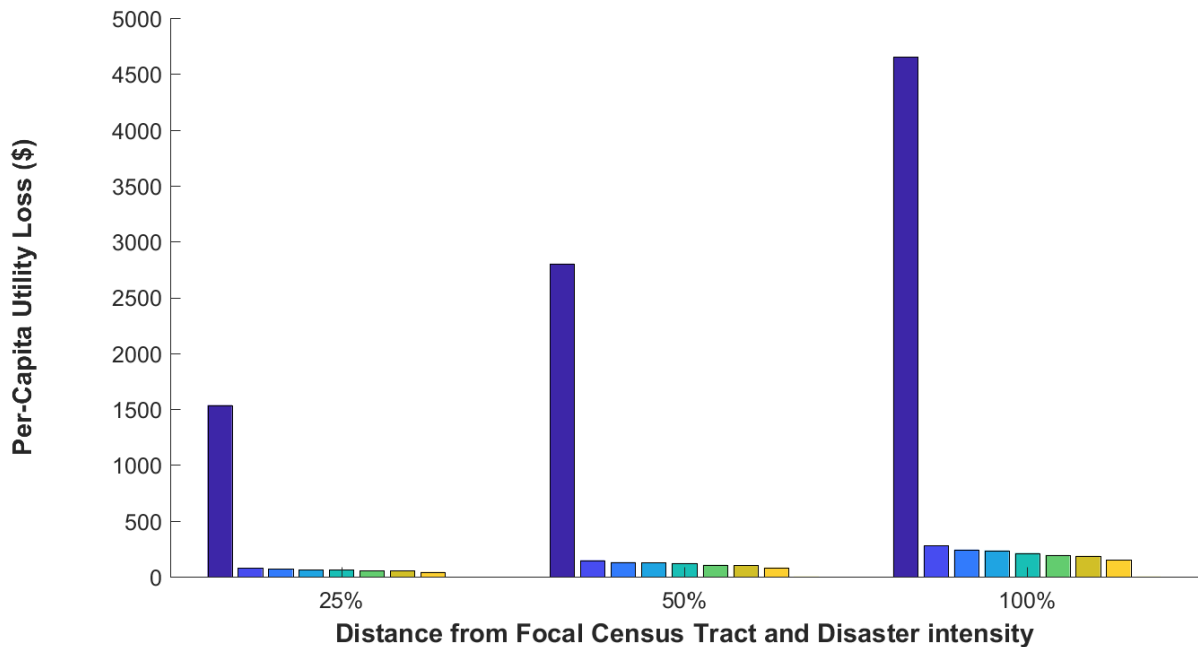


Figure 27: Expected Share of Total Utility Decreases Produced by a Natural Disaster Among Geographic Areas Defined by Distances from the Focal Tract, by Disaster Severity (25%/50%/100% of Jobs Lost)

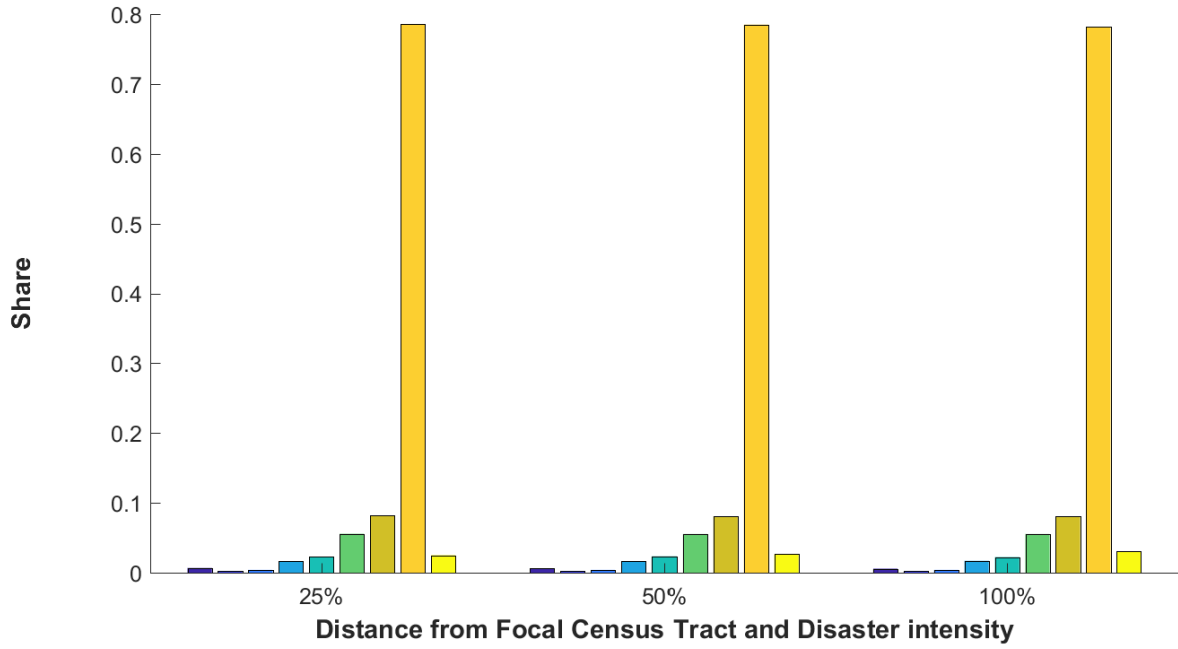


Figure 28: Expected Share of Additional Nonemployment Produced by a Natural Disaster among Workers Initially Employed (or Nonemployed) at Different Initial Earnings Quintiles (or Nonemployment), by Disaster Severity (25%/50%/100% of Jobs Lost)

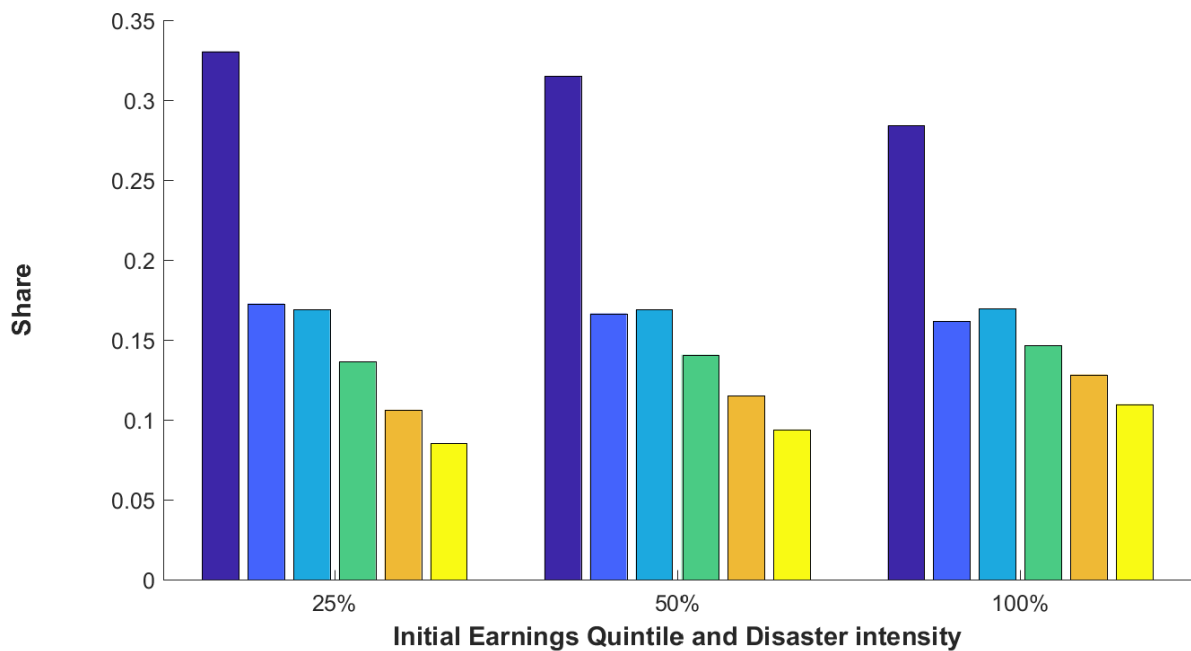


Figure 29: Change in P(Employed) Produced by a Natural Disaster (25% Jobs Lost) for Randomly Chosen Workers Initially Employed at Different Combinations of Initial Earnings Quintile (or Nonemployed) and Distance from Focal Tract

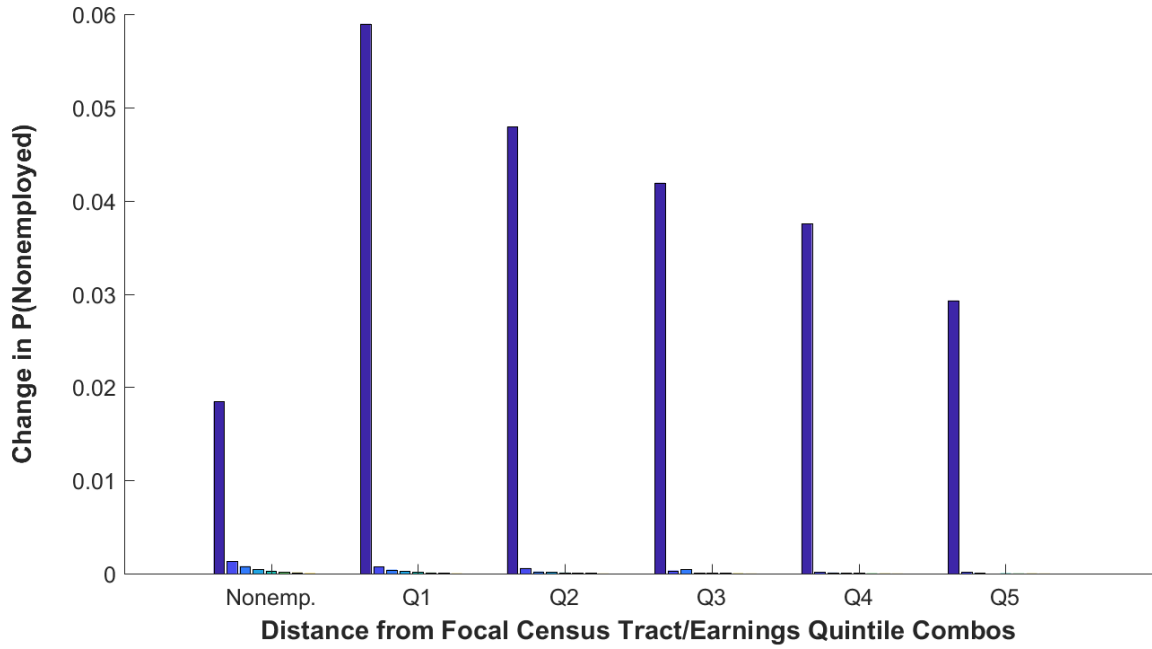


Figure 30: Change in P(Employed) Produced by a Natural Disaster (100% Jobs Lost) for Randomly Chosen Workers Initially Employed at Different Combinations of Initial Earnings Quintile (or Nonemployed) and Distance from Focal Tract

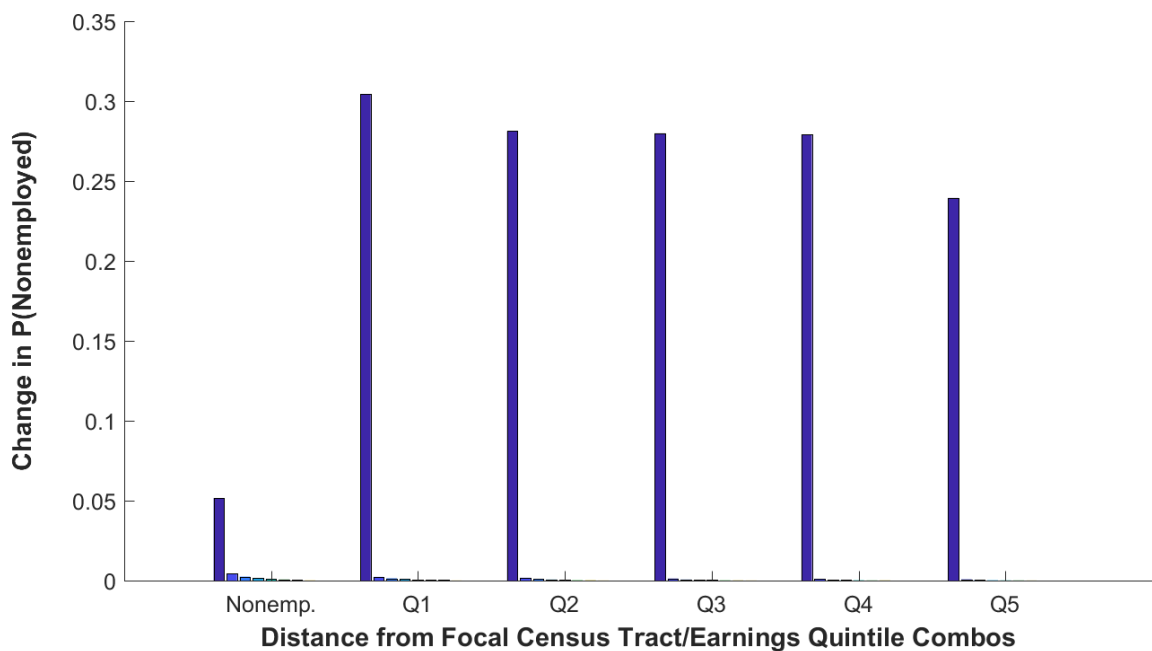


Figure 31: Expected Share of Earnings Losses Produced by a Natural Disaster among Workers Initially Employed (or Nonemployed) at Different Initial Earnings Quintiles (or Nonemployment), by Disaster Severity (25%/50%/100% of Jobs Lost)




Figure 32: Expected Decrease in Utility Produced by a Natural Disaster (25% Jobs Lost) for Randomly Chosen Workers Initially Employed at Different Combinations of Initial Earnings Quintile (or Nonemployed) and Distance from Focal Tract

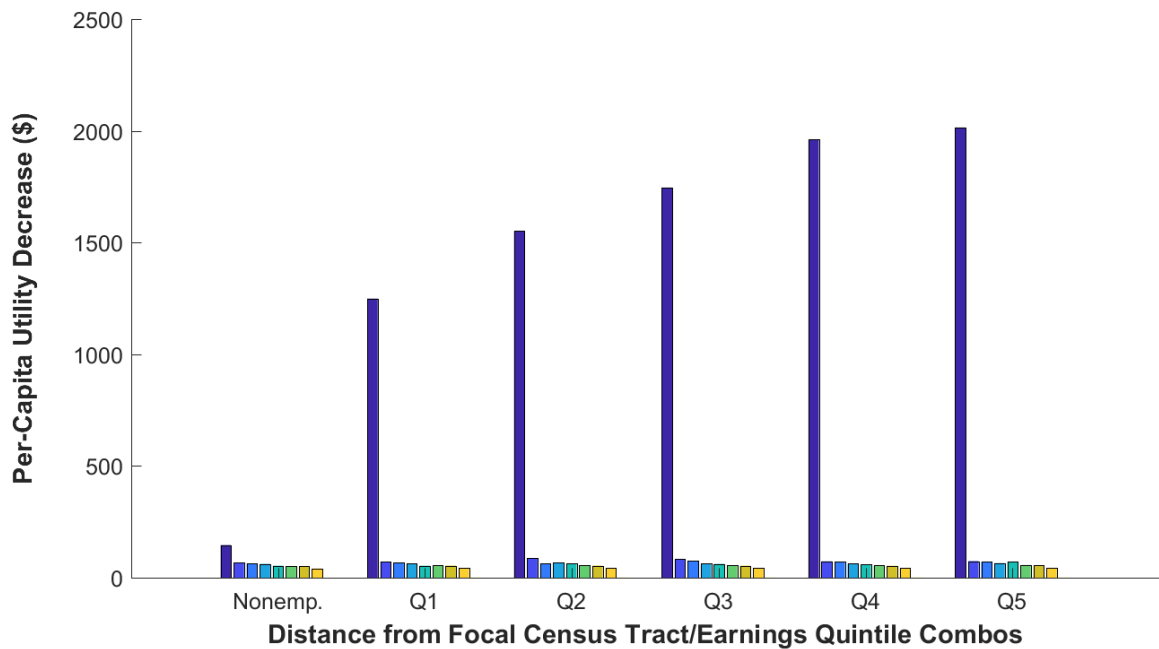


Figure 33: Expected Decrease in Utility Produced by a Natural Disaster (100% Jobs Lost) for Randomly Chosen Workers Initially Employed at Different Combinations of Initial Earnings Quintile (or Nonemployed) and Distance from Focal Tract

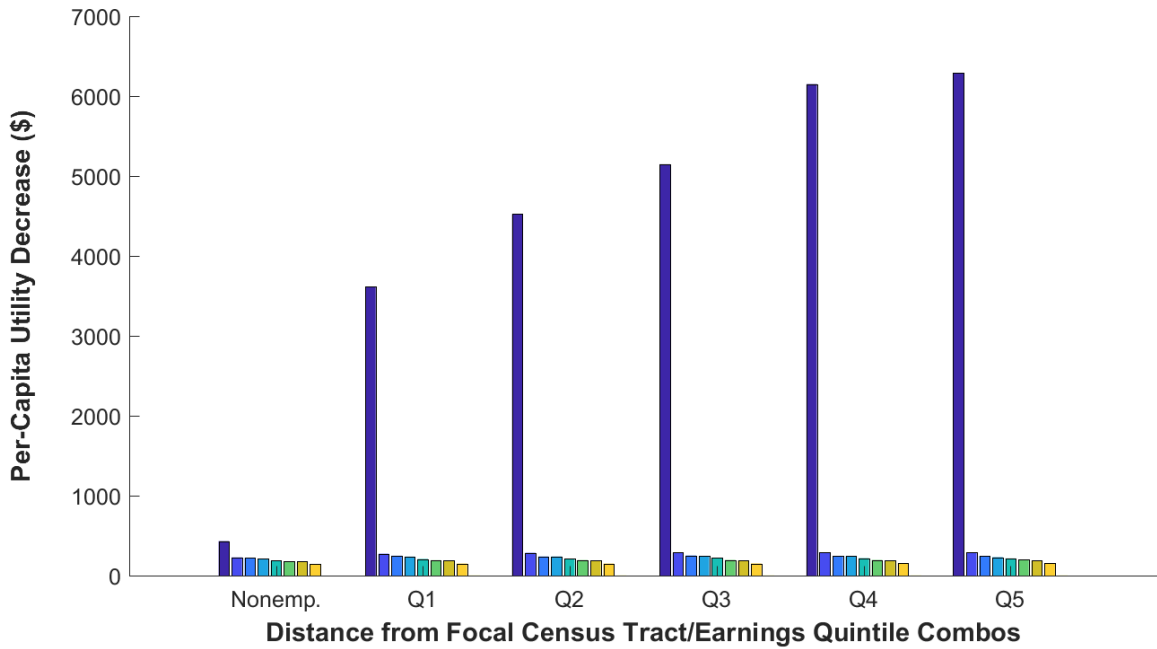


Figure 34: Change in Probability of Destination Employment (or Nonemployment) at Different Distances from Focal Tract after a Natural Disaster for Workers Initially Employed in the Focal Tract (Averaging Across the Initial Earnings Distribution), by Disaster Severity (25%/50%/100% of Jobs Lost)

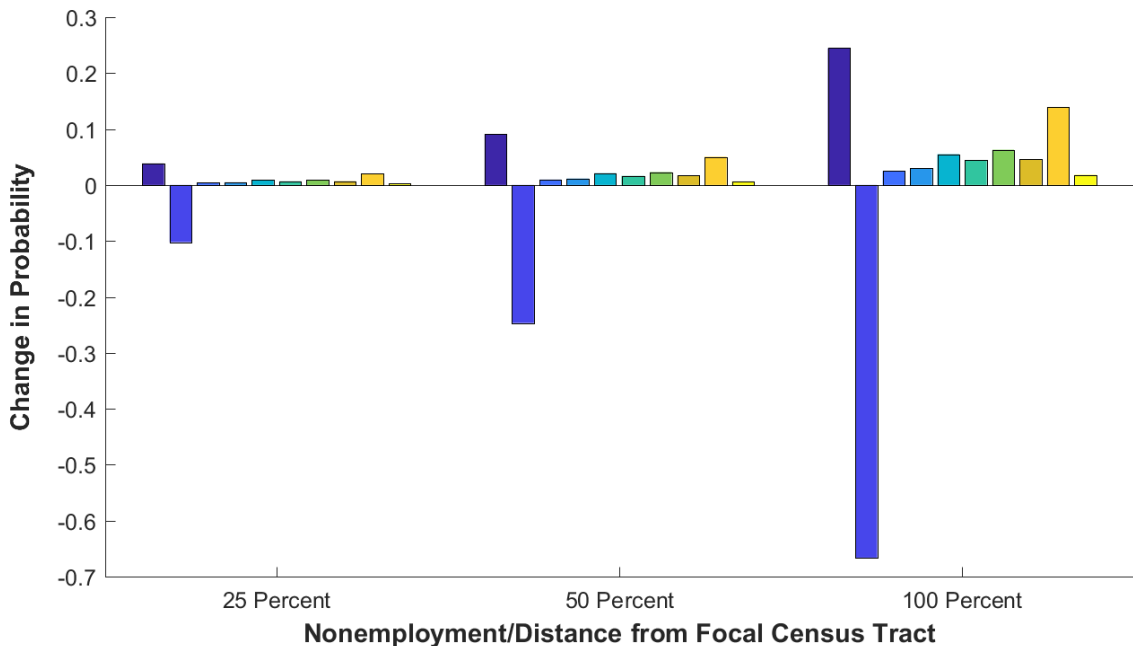


Figure 35: Change in Probability of Destination Employment (or Nonemployment) at Different Distances from the Focal Tract after a Natural Disaster (100% Jobs Lost) for Workers Initially Employed in the Focal Tract by Initial Earnings Quintile


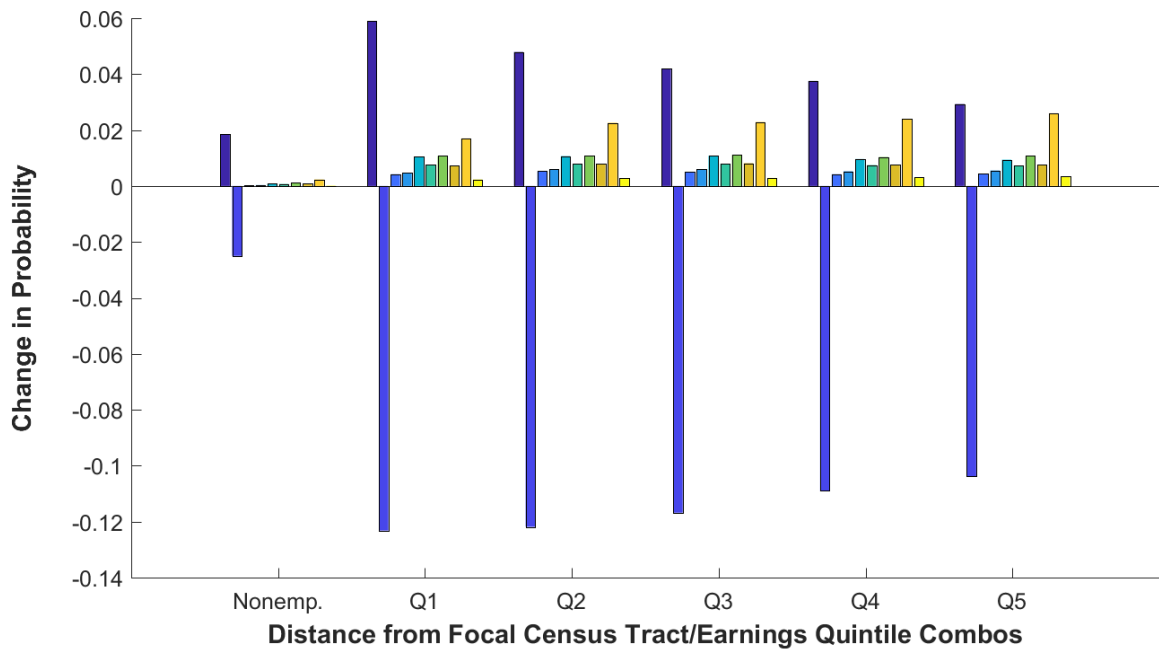


Figure 36: Change in Probability of Destination Employment (or Nonemployment) at Different Distances from the Focal Tract after a Natural Disaster (25% Jobs Lost) for Workers Initially Employed in the Focal Tract by Initial Earnings Quintile



Appendix

A1 Proof of Proposition A1

Proposition A1:

Suppose the following assumptions hold:

1') The assumptions laid out in sections 2.3 and 3 continue to hold. Namely, each joint surplus π_{ijk} is additively separable in the group-level and idiosyncratic components, the vector of idiosyncratic components $\epsilon_{ij(i)k}$ is independently and identically distributed, and follows the type 1 extreme value distribution, and Assumptions 1 and 2 hold.

2') The set of destination positions $k \in \tilde{\mathcal{K}}$ that will be filled in the stable counterfactual assignment are known in advance, and the set of destination positions $k \in \tilde{\mathcal{K}}$ that will remain unfilled in the stable counterfactual assignment are ignorable, in the sense that their existence does not change the assignment nor the division of surplus among the remaining set of positions \mathcal{K} and set of workers \mathcal{I} .

$$3') \frac{1}{|g_i|} \sum_{k:g(i,j(i),k)=g} e^{-\frac{q_k}{\sigma}} \approx \frac{1}{|d|} \sum_{k:d(k)=d(g)} e^{-\frac{q_k}{\sigma}} = C_{d(g)} \quad \forall (g, i).$$

$$4') P(g|i, d(g)) \approx P(g|o(g), d(g)) \quad \forall (g, i).$$

Then the group-level assignment $P^{CF}(g)$ that satisfies the following $O-1$ excess demand equations represents the unique group-level equilibrium assignment $P^{CF*}(g)$ consistent with the unique worker/firm level stable matching μ^{CF} :

$$\begin{aligned} \sum_{d \in \mathcal{D}} h^{CF}(d) \left(\sum_{g:o(g)=2} P^{CF}(g|d, C_2^{CF}, \dots, C_O^{CF}) \right) &= f^{CF}(2) \\ \vdots \\ \sum_{d \in \mathcal{D}} h^{CF}(d) \left(\sum_{g:o(g)=O} P^{CF}(g|d, C_2^{CF}, \dots, C_O^{CF}) \right) &= f^{CF}(O) \end{aligned} \quad (27)$$

where $P^{CF}(g|d, C_2^{CF}, \dots, C_O^{CF})$ is given by:

$$P^{CF}(g|d) = \frac{e^{\frac{\theta_g^{CF}}{\sigma}} P^{CF}(z(g)|o(g), d) f^{CF}(o(g)) C_o^{CF}}{\sum_{o' \in \mathcal{O}} \sum_{g' \in (o, d)} e^{\frac{\theta_{g'}^{CF}}{\sigma}} P^{CF}(z(g')|o'(g'), d) f^{CF}(o') C_{o'}^{CF}} \quad \forall d \in [1, \dots, D] \quad (28)$$

Proof: Proposition A1 states that assignment $P^{CF}(g)$ implied by the vector of mean utility values $\mathbf{C}^{CF} = [1, C_2, \dots, C_O^{CF}]$ that solves the system of equations (27) in fact represents the unique group-level stable (and equilibrium) assignment $P^{CF*}(g)$.

First, note that if unfilled positions are ignorable for the counterfactual assignment, then we can focus on finding a stable assignment of a restricted version of the assignment game in which only remaining K positions need to be considered. As discussed in footnote 20, Assumption 2' implicitly requires that no position that remains unfilled is ever the second-best option for any worker who takes a job in the destination period.

Furthermore, Assumption 2' imposes that each of the remaining positions will be filled in any stable matching. Recall that stability in the individual-level matching μ^{CF} requires:

$$\mu_{ij(i)k}^{CF} = 1 \text{ iff } k \in \arg \max_{k \in \tilde{K} \cup 0} \pi_{ij(i)k} - q_k^{CF} \text{ and } i \in \arg \max_{i \in \tilde{I} \cup 0} \pi_{ij(i)k} - r_i^{CF} \quad (29)$$

Assumption 2' allows us to replace $i \in \arg \max_{i \in \tilde{I} \cup 0} \pi_{ij(i)k} - r_i^{CF}$ with $i \in \arg \max_{i \in \tilde{I}} \pi_{ij(i)k} - r_i^{CF}$. In other words, we assume in advance that the individual rationality conditions that any proposed match yield a higher payoff to the firm than remaining vacant, $\pi_{ijk} - r_i > \pi_{0k}$ when $\mu_{ik} = 1$, are satisfied and can be ignored. Implicitly, this requires that the joint surpluses to workers and firms from matching up are sufficiently large relative to both workers' and firms' outside options.⁴¹ Imposing Assumption 2' will probably cause utility losses among local workers from negative local labor demand shocks to be overstated, since some workers would likely find jobs at positions that were not willing to hire at the original wage level but would enter the labor market at lower wage levels. Conversely, gains to local workers from positive shocks may be understated, since some local firms that filled positions at the original wage levels might choose to remain vacant (or move to other locations) when competition for local workers becomes more fierce.

In our applications the number of positions that will be filled (which Assumption 2' imposes will be known) is greater than the number of workers seeking positions (I). In order to be able to consistently allocate workers to transition groups, even when they move to (or remain in) nonemployment, we define a “nonemployment” destination type as the last destination type D . Because the number of workers who end up nonemployed is assumed to be known, we allocate enough “nonemployment” positions within type D , $h^{CF}(D)$, so that the number of workers I equals the number of “positions” K , once K includes the dummy nonemployment positions. We then normalize this common number of worker and firm positions (assumed to be very large) to be 1, and reinterpret $f^{CF}(o)$ and $h^{CF}(d)$ as probability mass functions providing shares of the relevant worker and position populations rather than counts.

As discussed in Section 3, Assumption 1', when combined with the stability conditions

⁴¹This implicitly requires that the unobserved draws ϵ_{0k} for firm vacancy values are taken from a bounded distribution rather than the Type 1 extreme value distribution.

(29), implies that the probability that a given position k will be filled by a particular worker i is given by the logit form (11). When combined with Assumptions 1 and 2 (also cited by Assumption 1'), this implies that the group-level conditional choice probability $P(g|d)$ takes the form (28) for any destination types d that are composed of positions k (as derived in section 3).

However, note the statement of Proposition A1 makes it clear that the form (28) also holds for the last type D , which contains the “dummy” nonemployment positions whose “choices” will be workers moving to nonemployment. The stability conditions (29) do not provide any justification for why these dummy nonemployment positions should be filled via the same logit form as the other destination types that consist of actual positions at firms. Thus, the inclusion of these dummy positions, and the assumption that the probability distribution over alternative groups representing different worker and job match characteristics $(o(g), z(g))$ follows the logit form, are mere computational devices to calculate the equilibrium assignment. That this computational device in fact yields the unique stable assignment for the proposed counterfactual labor market is the primary reason Proposition A1 requires a proof.

However, the stability conditions and Assumption 1' imply that the probability that a given worker i will choose a particular firm k (where $k = 0$ represents nonemployment) is also given by the logit form (Decker et al. (2013)):

$$P^{CF}(k|i) = \frac{e^{\frac{\theta_g^{CF} - q_k^{CF}}{\sigma}}}{\sum_{k' \in \mathcal{K} \cup 0} e^{\frac{\theta_{g'}^{CF} - q_{k'}^{CF}}{\sigma}}} \quad (30)$$

This can then be aggregated (using the same steps as in section 3) to provide an expression for the probability that a randomly chosen worker from a given origin type o matches with a position that yields a transition assigned to group g :

$$P^{CF}(g|o) = \frac{1}{|o|} \sum_{i \in o} \frac{(e^{\frac{\theta_g^{CF}}{\sigma}}) (\sum_{k: g(i, j(i), k) = g} e^{\frac{-q_k^{CF}}{\sigma}})}{\sum_{k' \in \mathcal{K} \cup 0} e^{\frac{\theta_{g'}^{CF} - q_{k'}^{CF}}{\sigma}}} \quad (31)$$

Assumptions 3' and 4', which are analogues to Assumptions 1 and 2 in section 3, allow us to simply this expression to the following:

$$P^{CF}(g|o) = \frac{e^{\frac{\theta_g^{CF}}{\sigma}} P^{CF}(z(g)|o, d(g)) h^{CF}(d(g)) \tilde{C}_d^{CF}}{\sum_{d' \in \mathcal{D}} \sum_{g' \in (o, d')} e^{\frac{\theta_{g'}^{CF}}{\sigma}} P^{CF}(z(g')|o'(g'), d) h^{CF}(d') \tilde{C}_{d'}^{CF}} \quad \forall o \in [1, \dots, O] \quad (32)$$

Assumption 3' states that the discounted profits of alternative positions k assigned to the same destination type d are roughly the same. This implies that the profit share that workers

must provide to the position in a stable matching is approximately the same for their existing positions as for other positions in the same local area featuring the same industry and firm size and firm average pay categories, and can be summarized by a parameter C_d^{CF} that is defined at the destination-type level.

Taken literally (given the characteristics we use to define groups), Assumption 4' states that every worker assigned to the same origin type starts the year in firms with the same number of destination positions, which clearly does not hold. More broadly, though, Assumptions 3' and 4' allow us to replace the term $\sum_{k:g(i,j(i),k)=g} e^{\frac{-q_k^{CF}}{\sigma}}$ that depends on the individual i with an expression $P^{CF}(g|o, d(g))h^{CF}(d(g))\tilde{C}_{d(g)}^{CF}$ that depends on only group and destination-type level terms. Essentially, we are assuming that ignoring within-origin type variation in the number of positions at which they would be stayers (due to different firm sizes of initial job matches) when aggregating is not generating significant bias in the counterfactual assignment and incidence estimates.⁴²

Under Assumptions 1' through 4', the group-level stable matching must satisfy the following market clearing conditions, which specify that supply must equal demand for each destination position type d :

$$\sum_{o \in \mathcal{O}} f^{CF}(o) \left(\sum_{g:d(g)=2} P^{CF*}(g|o, \tilde{\mathbf{C}}^{\mathbf{CF}}) \right) = h^{CF}(2) \quad (33)$$

$$\vdots \quad (34)$$

$$\sum_{o \in \mathcal{O}} f^{CF}(o) \left(\sum_{g:d(g)=D} P^{CF*}(g|o, \tilde{\mathbf{C}}^{\mathbf{CF}}) \right) = h^{CF}(D) \quad (35)$$

where $\tilde{\mathbf{C}}^{\mathbf{CF}}$ represents the $D - 1$ length vector $= [1, \tilde{C}_2^{CF}, \dots, \tilde{C}_D^{CF}]$ and each conditional probability $P^{CF*}(g|o, \tilde{\mathbf{C}}^{\mathbf{CF}})$ takes the form in (32).

Note in particular that Assumption 2' allows us to ignore the possibility that supply might exceed demand for some destination position types (implying some vacant positions). In this alternative position-side system of equations, the expressions for each conditional probability $P^{CF*}(g|o)$ do in fact stem directly from the necessary stability conditions. And all of the feasibility conditions for a stable matching are incorporated into the zero-excess demand equations (since $P^{CF*}(g|o)$ sum to 1 by construction, the assignment $P^{CF*}(g)$ that satisfies this system necessarily sums to the origin-type PMF $f^{CF}(o)$). Thus, the proof by Decker et al. (2013) that there exists a unique group-level assignment that satisfies all of the group-level feasibility and stability conditions (and is thus consistent with a stable matching in the assignment game defined at the level of worker-position matches) applies here.

If one wished, one could directly compute the unique group-level counterfactual assign-

⁴²Note also that assumptions 3' and 4' may not be necessary conditions for Proposition A1. We are seeking an alternative proof of Proposition A1 that does not require assumptions 3' and 4'.

ment $P^{CF*}(g|o)$ by finding a $D - 1$ length vector $\tilde{\mathbf{C}}^{\mathbf{CF}}$ that solved this system, and constructing the implied assignment by plugging this vector into the conditional probability expressions (32). However, when $D \gg O$, solving this system is considerably more computationally burdensome than solving the worker-side counterpart (27), which features $O - 1$ equations. Thus, the remainder of this proof is devoted to showing that any assignment $P^{CF}(g)$ implied by a solution to (27) must equal the assignment $P^{CF*}(g)$ implied by a solution to (35). And since we know that the latter solution represents the unique group-level matching consistent with stability in the assignment game, the former solution must also be unique, and must also represent the group-level matching consistent with stability in the assignment game. Essentially, this amounts to showing that the device of adding “dummy” nonemployment positions present in (27) appropriately incorporates the surpluses π_{i0} that workers obtain from staying single.

Consider an O length vector $\mathbf{C}^{\mathbf{CF}} = [1, C_2^{CF}, \dots, C_O^{CF}]$ that solves (27) and generates assignment $P^{CF}(g)$. We wish to show that one can use $\mathbf{C}^{\mathbf{CF}}$ to construct an alternative D length vector $\tilde{\mathbf{C}}^{\mathbf{CF}} = [1, \tilde{C}_2^{CF}, \dots, \tilde{C}_D^{CF}]$ that solves (35), and that the assignment it generates, $P^{CF*}(g)$, equals $P^{CF}(g)$.

We propose the following vector $\tilde{\mathbf{C}}^{\mathbf{CF}}$:

$$\tilde{C}_d^{CF} = \frac{\sum_{o=1}^O \sum_{g':(o(g'),d(g'))=(o,D)} e^{\frac{\theta_{g'}}{\sigma}} f^{CF}(o) P(z(g')|o, D) C_o^{CF}}{\sum_{o=1}^O \sum_{g':(o(g'),d(g'))=(o,d)} e^{\frac{\theta_{g'}}{\sigma}} f^{CF}(o) P(z(g')|o, d) C_o^{CF}} \quad \forall d \in [1, \dots, D] \quad (36)$$

Here, the numerator captures the inclusive value (as defined by Menzel (2015)) associated with the nonemployment destination type D , while the denominator captures the inclusive value for the chosen destination type d . This implies that $\tilde{C}_D^{CF} = 1$. While any destination type could be chosen as the one whose mean exponentiated profit value is normalized, normalizing the nonemployment type is particularly appealing, since it implies “profit” values of 0 for the dummy nonemployment destination type D ($\tilde{C}_D^{CF} = e^{\bar{q}_D} = e^0 = 1$).

To conserve notation, let λ represent the inclusive value associated with the nonemployment destination type D , the numerator in (41):

$$\lambda = \sum_{o=1}^O \sum_{g':(o(g'),d(g'))=(o,D)} e^{\frac{\theta_{g'}}{\sigma}} f^{CF}(o) P^{CF}(z(g')|o, D) C_o^{CF} \quad (37)$$

Note that λ is independent of destination type.

We begin by showing that the assignments implied by the vectors $[C_1^{CF}, \dots, C_O^{CF}]$ and $[C_1^{CF}, \dots, \tilde{C}_D^{CF}]$ are identical: $P^{CF}(g) = P^{CF*}(g)$.

Note first that since \mathbf{C}^{CF} solves the worker-side system of excess demand equations (27),

we know that

$$\begin{aligned}
& \sum_{d' \in \mathcal{D}} h^{CF}(d') \sum_{g' \in (o, d')} \frac{e^{\frac{\theta_{g'}^{CF}}{\sigma}} P^{CF}(z(g')|o, d') f^{CF}(o) C_o^{CF}}{\sum_{o'=1}^O \sum_{g':(o(g'), d(g'))=(o', d)} e^{\frac{\theta_{g'}}{\sigma}} f^{CF}(o') P(z(g')|o', d) C_{o'}^{CF}} = f^{CF}(o) \forall o \in [1, O] \\
& \Rightarrow \sum_{d' \in \mathcal{D}} \sum_{g' \in (o, d')} \frac{e^{\frac{\theta_{g'}^{CF}}{\sigma}} P^{CF}(z(g')|o, d') h^{CF}(d')}{\sum_{o=1}^O \sum_{g':(o(g'), d(g'))=(o, d)} e^{\frac{\theta_{g'}}{\sigma}} f^{CF}(o) P(z(g')|o, d) C_o^{CF}} = \frac{1}{C_o^{CF}} \forall o \in [1, O] \\
& \Rightarrow \sum_{d' \in \mathcal{D}} \sum_{g' \in (o, d')} \frac{e^{\frac{\theta_{g'}^{CF}}{\sigma}} P^{CF}(z(g')|o, d') h^{CF}(d')}{\frac{\lambda}{\tilde{C}_{d'}^{CF}}} = \frac{1}{C_o^{CF}} \forall o \in [1, O] \\
& \Rightarrow \sum_{d' \in \mathcal{D}} \sum_{g' \in (o, d')} e^{\frac{\theta_{g'}^{CF}}{\sigma}} P^{CF}(z(g')|o, d') h^{CF}(d') \tilde{C}_{d'}^{CF} = \frac{\lambda}{C_o^{CF}} \forall o \in [1, O] \tag{38}
\end{aligned}$$

We can now proceed:

$$\begin{aligned}
P^{CF*}(g) &= f^{CF}(o) P^{CF*}(g|o) = f^{CF}(o) \frac{e^{\frac{\theta_g^{CF}}{\sigma}} P^{CF}(z(g)|o, d) h^{CF}(d) \tilde{C}_d^{CF}}{\sum_{d' \in \mathcal{D}} \sum_{g' \in (o, d')} e^{\frac{\theta_{g'}}{\sigma}} P^{CF}(z(g')|o, d') h^{CF}(d') \tilde{C}_{d'}^{CF}} \\
&= \frac{f^{CF}(o) e^{\frac{\theta_g^{CF}}{\sigma}} P^{CF}(z(g)|o, d) h^{CF}(d) \tilde{C}_d^{CF} C_o^{CF}}{\lambda} \\
&= h^{CF}(d) \frac{e^{\frac{\theta_g^{CF}}{\sigma}} f^{CF}(o) P^{CF}(z(g)|o, d) \lambda C_o^{CF}}{\lambda \sum_{o'=1}^O \sum_{g':(o(g'), d(g'))=(o', d)} e^{\frac{\theta_{g'}}{\sigma}} f^{CF}(o') P(z(g')|o', d) C_{o'}^{CF}} \\
&= h^{CF}(d) P^{CF}(g|d) = P^{CF}(g) \tag{39}
\end{aligned}$$

It remains to show that the chosen $\tilde{\mathbf{C}}^{CF}$ vector (41) solves (35). Consider the left-hand side of the excess demand equation for an arbitrarily chosen destination type d in the system (35). We can write:

$$\begin{aligned}
& \sum_{o=1}^O \sum_{g:(o(g), d(g))=(o, d)} f^{CF}(o) P^{CF*}(g|o, \Theta^{CF}, \tilde{\mathbf{C}}^{CF}) \\
& \sum_{o=1}^O \sum_{g:(o(g), d(g))=(o, d)} h^{CF}(d) P^{CF}(g|d, \Theta^{CF}, \mathbf{C}^{CF}) \\
& h^{CF}(d) \sum_{o=1}^O \sum_{g:(o(g), d(g))=(o, d)} P^{CF}(g|d, \Theta^{CF}, \mathbf{C}^{CF}) \\
& = h^{CF}(d) \sum_{g:d(g)=d} P^{CF}(g|d, \Theta^{CF}, \mathbf{C}^{CF}) \\
& = h^{CF}(d) \tag{40}
\end{aligned}$$

where the last line uses the fact that $P^{CF}(g|d)$ is a (conditional) probability distribution and thus sums to one. Since we have proved that the implied “demand” by workers for positions of an arbitrary destination type equals the “supply” $h^{CF}(d)$, we have thus proved that \tilde{C}^{CF} solves the system (35).

Notice that the expression for the proposed equilibrium mean ex post profit vector (41) has value beyond its use in proving proposition A1. Once the O -vector of mean ex post utilities $\{C_o^{CF}\}$ for each origin type have been computed, we can use (41) to directly calculate the mean ex post profit vector for each destination position type d without having to solve a system of $D - 1$ equations. This is quite valuable when $D \gg O$, as it is in our application. Of course, the equivalent mapping can be inferred by symmetry for the opposite case where $O \gg D$:

$$C_o^{CF} = \frac{\sum_{d=1}^D \sum_{g':(o(g'),d(g'))=(O,d)} e^{\frac{\theta_{g'}}{\sigma}} h^{CF}(d) P(z = z(g')|O, d) \tilde{C}_d^{CF}}{\sum_{d=1}^D \sum_{g':(o(g'),d(g'))=(o,d)} e^{\frac{\theta_{g'}}{\sigma}} h^{CF}(d) P(z = z(g')|o, d) \tilde{C}_d^{CF}} \quad \forall o \in [1, \dots, O] \quad (41)$$

In section 3.2 we showed that these vectors are sufficient to determine both the worker and firm type-level incidence of any counterfactual shocks to the composition or spatial distribution of labor supply and/or labor demand. Thus, at least in cases where the proposed model is a reasonable approximation of the functioning of the labor market (and housing supply is sufficiently elastic and agglomeration effects and other product market spillovers are second order), a proper welfare analysis of such shocks only requires solving at most $\min\{O, D\}$ non-linear excess demand equations. Since an analytical Jacobian can be derived and fed as an input to non-linear equations solvers, relatively large scale assignment problems featuring thousands of types on one side of the market (and perhaps more on the opposite side) can be solved within a matter of minutes.

A2 Proof of Proposition 1

Proposition 1:

Define the set $\Theta^{D-in-D} \equiv \left\{ \frac{(\theta_g - \theta_{g'}) - (\theta_{g''} - \theta_{g'''})}{\sigma} \quad \forall (g, g', g'', g''') : o(g) = o(g''), o(g') = o(g'''), d(g) = d(g'), d(g'') = d(g''') \right\}$. Given knowledge of Θ^{D-in-D} , a set $\tilde{\Theta} = \{\tilde{\theta}_g\}$ can be constructed such that the unique group level assignment $P^{CF}(g)$ that satisfies the system of excess demand equations (26) using $\theta_g^{CF} = \tilde{\theta}_g \quad \forall g$ and arbitrary marginal distributions for origin and destination types $f^{CF}(\ast)$ and $g^{CF}(\ast)$ will also satisfy the corresponding system of excess demand equations using $\theta_g^{CF} = \theta_g \quad \forall g$ and arbitrary distributions $f^{CF}(\ast)$ and $g^{CF}(\ast)$. Furthermore, denote by $\{\tilde{C}_1^{CF}, \dots, \tilde{C}_O^{CF}\}$ the market-clearing utility values that clear the market using $\theta_g^{CF} = \tilde{\theta}_g$, and denote by $\{C_1^{CF}, \dots, C_O^{CF}\}$ the market-clearing utility values that clear the market using $\theta_g^{CF} = \theta_g$. Then $\{\tilde{C}_o^{CF}\}$ will satisfy $\tilde{C}_o^{CF} = C_o^{CF} e^{\frac{\Delta_o}{\sigma}} \quad \forall o$

for some set of origin type-specific constants $\{\Delta_o\}$ that is invariant to the choice of $f^{CF}(\ast)$ and $g^{CF}(\ast)$.

Proof: We prove Proposition 1 by construction.

Let $z(i, j, k) = 1(m(j) = m(k))$ represent an indicator that takes on the value of 1 if the firms associated with positions j and k are the same, and 0 otherwise. Recall also that all worker transitions assigned to the same transition group g share values of the worker and firm characteristics that define the worker's origin and firm's destination types o and d , respectively, as well as the value of the indicator $z(i, j, k)$. Thus, we can write $o(g)$, $d(g)$ and $z(g)$ for any group g . Let the origin types be ordered (arbitrarily) from $o = 1 \dots o = O$, and let the destination types be ordered (arbitrarily) from $d = 1 \dots d = D$. Let $g(o, d, z)$ denote the group associated with origin type o , destination type d , and existing worker indicator z . Assume that the set $\Theta^{D-in-D} = \left\{ \frac{(\theta_g - \theta_{g'}) - (\theta_{g''} - \theta_{g'''})}{\sigma} \forall (g, g', g'', g''') \right\}$ is known, since a consistent estimator for each element of the set can be obtained via adjusted log odds ratios, as described in Section 3.

Consider defining the following set of alternative group-level joint surplus values $\tilde{\Theta} = \{\tilde{\theta}_g\}$ as follows:

$$\tilde{\theta}_{g'} = 0 \forall g' : (o(g') = 1 \text{ and/or } d(g') = 1) \text{ and } z(g') = 0 \quad (42)$$

$$\tilde{\theta}_g = \frac{(\theta_{g'} - \theta_{g(1,d(g'),0)}) - (\theta_{g(o(g'),1,0)} - \theta_{g(1,1,0)})}{\sigma} \forall g' : (d(g') \neq \text{ and } o(g') \neq 1) \text{ and/or } z(g') \neq 0 \quad (43)$$

Under the definitions in (42) and (43), we have:

$$\begin{aligned} \frac{(\tilde{\theta}_g - \tilde{\theta}_{g'}) - (\tilde{\theta}_{g''} - \tilde{\theta}_{g'''})}{\sigma} &= \frac{(\theta_g - \theta_{g'}) - (\theta_{g''} - \theta_{g'''})}{\sigma} \\ \forall (g, g', g'', g''') : o(g) &= o(g''), o(g') = o(g'''), d(g) = d(g'), d(g'') = d(g''') \end{aligned} \quad (44)$$

Thus, the appropriate difference-in-differences using elements of $\tilde{\Theta}$ match the corresponding difference-in-differences of true surpluses in Θ^{D-in-D} , so that all of the information about Θ contained in the identified set Θ^{D-in-D} is retained. Furthermore, unlike the true set Θ , the construction of $\tilde{\Theta}$ only requires knowledge of Θ^{D-in-D} .

Next, note that the elements of $\tilde{\Theta}$ can be written in the following form:

$$\tilde{\theta}_g = \theta_g + \Delta_{o(g)}^1 + \Delta_{d(g)}^2 \forall g \in \mathcal{G}, \text{ where} \quad (45)$$

$$\Delta_{o(g)}^1 = \theta_{g(o(g),1,0)} - \theta_{g(1,1,0)} \quad (46)$$

$$\Delta_{d(g)}^2 = \theta_{g(1,d(g),0)} \quad (47)$$

where \mathcal{G} is the set of all possible transition groups. In other words, each alternative surplus

$\tilde{\theta}_g$ equals the true surplus θ_g plus a constant ($\Delta_{o(g)}^1$) that is common to all groups featuring the same origin type and a constant ($\Delta_{d(g)}^2$) that is common to all groups featuring the same destination type.

Next, recall that there exists a unique aggregate assignment associated with each combination of marginal origin and destination type distributions $f^{CF}(o)$ and $h^{CF}(d)$ and set of group-level surpluses, including $\tilde{\Theta}$. Let $\tilde{P}^{CF}(\ast) \equiv P^{CF}(\ast|\tilde{\Theta}, \tilde{C}_2^{CF}, \dots, \tilde{C}_O^{CF})$ represent the unique counterfactual assignment that results from combining arbitrary marginal distributions $f^{CF}(o)$ and $h^{CF}(d)$ with the set $\tilde{\Theta}$. $\tilde{\mathbf{C}}^{\mathbf{CF}} = [1, \tilde{C}_2^{CF} \dots \tilde{C}_O^{CF}]$ denotes the vector of mean exponentiated utility values for each origin type o (with \tilde{C}_1^{CF} normalized to 1) that solves the system of excess demand equations below, and thus yields $\tilde{P}^{CF}(g) \forall g \in \mathcal{G}$ when plugged into equation (25) along with the elements of $\tilde{\Theta}$:

$$\begin{aligned} \sum_{d \in \mathcal{D}} h^{CF}(d) \left(\sum_{g: o(g)=2} P^{CF}(g|d, \tilde{\Theta}, \tilde{\mathbf{C}}^{\mathbf{CF}}) \right) &= f^{CF}(2) \\ \vdots \\ \sum_{d \in \mathcal{D}} h^{CF}(d) \left(\sum_{g: o(g)=O} P^{CF}(g|d, \tilde{\Theta}, \tilde{\mathbf{C}}^{\mathbf{CF}}) \right) &= f^{CF}(O) \end{aligned} \quad (48)$$

We wish to show that $\tilde{P}^{CF}(\ast) \equiv P^{CF}(\ast|\tilde{\Theta}, \tilde{\mathbf{C}}^{\mathbf{CF}})$ will be identical to the alternative unique counterfactual equilibrium assignment $P^{CF}(\ast|\Theta, \mathbf{C}^{\mathbf{CF}})$ that combines the same arbitrary marginal distributions $f^{CF}(o)$ and $h^{CF}(d)$ with the set Θ instead of $\tilde{\Theta}$. Here, $\mathbf{C}^{\mathbf{CF}} = [1, C_2^{CF} \dots C_O^{CF}]$ denotes a vector of o -type-specific mean exponentiated utility values that clears the market by satisfying the following alternative excess demand equations:⁴³

$$\begin{aligned} \sum_{d \in \mathcal{D}} h^{CF}(d) \left(\sum_{g: o(g)=2} P^{CF}(g|d, \Theta, \mathbf{C}^{\mathbf{CF}}) \right) &= f^{CF}(2) \\ \vdots \\ \sum_{d \in \mathcal{D}} h^{CF}(d) \left(\sum_{g: o(g)=O} P^{CF}(g|d, \Theta, \mathbf{C}^{\mathbf{CF}}) \right) &= f^{CF}(O) \end{aligned} \quad (49)$$

Since all other terms are unchanged between the systems (48) and (49), it suffices to show that $P^{CF}(g|d, \tilde{\Theta}, \tilde{\mathbf{C}}^{\mathbf{CF}}) = P^{CF}(g|d, \Theta, \mathbf{C}^{\mathbf{CF}}) \forall g \in \mathcal{G}$ for some vector $\mathbf{C}^{\mathbf{CF}}$. Consider the following proposed vector $\mathbf{C}^{\mathbf{CF}}$:

$$C_o^{CF} = \tilde{C}_o^{CF} e^{\frac{\Delta_o^1}{\sigma}} \quad \forall o \in [2, \dots, O] \quad (50)$$

⁴³Note that we have suppressed the dependence of $P^{CF}(\ast|\Theta, \mathbf{C}^{\mathbf{CF}}, f^{CF}(o), h^{CF}(d), P(z = z(g)|o, d))$ on $f^{CF}(o)$, $h^{CF}(d)$, and $P(z(g)|o, d)$ because these are held fixed across the two alternative counterfactual simulations.

where Δ_o^1 is as defined in (46). For an arbitrary choice of g , we obtain:

$$\begin{aligned}
& P^{CF}(g|d(g), \tilde{\Theta}, \tilde{\mathbf{C}}^{CF}) \\
&= \frac{e^{\frac{\tilde{\theta}_g^{CF}}{\sigma}} P^{CF}(z(g)|o(g), d(g)) f^{CF}(o(g)) \tilde{C}_o^{CF}}{\sum_{o' \in \mathcal{O}} \sum_{g' \in (o, d)} e^{\frac{\tilde{\theta}_{g'}^{CF}}{\sigma}} P^{CF}(z(g')|o'(g'), d(g)) f^{CF}(o') \tilde{C}_{o'}^{CF}} \\
&= \frac{e^{\frac{(\theta_g^{CF} + \Delta_o^1 + \Delta_d^2)}{\sigma}} P^{CF}(z(g)|o(g), d(g)) f^{CF}(o(g)) C_o^{CF} e^{-\frac{\Delta_o^1}{\sigma}}}{\sum_{o' \in \mathcal{O}} \sum_{g' \in (o', d)} e^{\frac{(\theta_{g'}^{CF} + \Delta_{o'}^1 + \Delta_{d'}^2)}{\sigma}} P^{CF}(z(g')|o'(g'), d(g)) f^{CF}(o') C_{o'}^{CF} e^{-\frac{\Delta_{o'}^1}{\sigma}}} \\
&= e^{\frac{\Delta_o^1}{\sigma}} e^{\frac{\Delta_d^2}{\sigma}} e^{-\frac{\Delta_o^1}{\sigma}} \frac{e^{\frac{\theta_g^{CF}}{\sigma}} P^{CF}(z(g)|o(g), d(g)) f^{CF}(o(g)) C_o^{CF}}{e^{\frac{\Delta_d^2}{\sigma}} \sum_{o' \in \mathcal{O}} e^{\frac{\Delta_{o'}^1}{\sigma}} e^{-\frac{\Delta_{o'}^1}{\sigma}} \sum_{g' \in (o', d)} e^{\frac{\theta_{g'}^{CF}}{\sigma}} P^{CF}(z(g')|o'(g'), d(g)) f^{CF}(o') C_{o'}^{CF}}} \\
&= \frac{e^{\frac{\theta_g^{CF}}{\sigma}} P^{CF}(z(g)|o(g), d(g)) f^{CF}(o(g)) C_o^{CF}}{\sum_{o' \in \mathcal{O}} \sum_{g' \in (o', d)} e^{\frac{\theta_{g'}^{CF}}{\sigma}} P^{CF}(z(g')|o'(g'), d(g)) f^{CF}(o') C_{o'}^{CF}} \\
&= P^{CF}(g|d, \Theta, \mathbf{C}^{CF}) \tag{51}
\end{aligned}$$

This proves that $P^{CF}(g|d, \Theta, \mathbf{C}^{CF})$ also satisfies the market clearing conditions (49) above, and will therefore be the unique group-level assignment consistent with marketwide equilibrium and stability. Thus, we have shown that the counterfactual assignment we recover when using an alternative set of surpluses $\tilde{\Theta}$ derived from the identified set Θ^{D-in-D} will in fact equal the counterfactual assignment we desire, which is based on the true set of joint surplus values Θ . Furthermore, while origin-type specific mean utility values $\tilde{\mathbf{C}}^{CF}$ that clears the market given $\tilde{\Theta}$ will differ for each origin type from the corresponding vector \mathbf{C}^{CF} based on the true surplus set Θ , these differences are invariant to the marginal origin and destination distributions $f^{CF}(o)$ and $h^{CF}(d)$ used to define the counterfactual. This implies that differences in utility gains caused by alternative counterfactuals among origin groups are identified, permitting comparisons of the utility incidence of alternative labor supply or demand shocks. This concludes the proof.

A3 Estimating the Value of σ

We attempt to estimate σ , the standard deviation of the unobserved match-level component $\epsilon_{ij(i)k}$, by exploiting the fact that the composition of U.S. origin and destination job matches $f^y(o)$ and $h^y(d)$ evolved across years y . Specifically, we estimate the set of group-level surpluses $\{\theta_g^{2007}\}$ from the observed 2007-2008 matching. Then, holding these surplus values fixed, we combine $\{\theta_g^{2007}\}$ with $f^y(o)$ and $h^y(d)$ from each other year $y \in [1993, 2010]$ to generate counterfactual assignments and changes in scaled mean (exponentiated) utility values $\{C_o^{CF}\}$ for each origin type. These counterfactuals predict how mean worker utilities

by skill/location combination could have been expected to evolve over the observed period given the observed compositional changes in labor supply and demand had the underlying surplus values $\{\theta_g\}$ been constant and equal to $\{\theta_g^{2007}\}$ throughout the period.

To the extent that most of evolution in the utility premia enjoyed by workers in particular locations and skill categories was due primarily to changes in supply and demand composition rather than changes in the moving costs, recruiting costs, tastes, and relative productivities that compose the joint surplus values $\{\theta_g\}$, these counterfactual predictions will be reasonable approximations of the realized evolution of ex post utility over time by origin type. Recall that $C_o^{CF} \approx \frac{1}{|o|} \sum_{i:o(i,j(i))=o} e^{-\frac{r_i^{CF}}{\sigma}}$. Thus, if ex post utility r_i^{CF} does not vary too much across individuals within an origin type, so that Jensen's inequality is near equality and $\frac{1}{|o^y|} \sum_{i:o(i,j(i))=o} e^{-\frac{r_i^{CF,y}}{\sigma^y}} \approx e^{-\frac{\bar{r}_o^{CF,y}}{\sigma^y}}$, then taking logs yields $\ln(C_o^{CF,y}) \approx \frac{\bar{r}_o^{CF,y}}{\sigma^y}$.

Next, we form the corresponding changes in observed annual earnings from origin to destination match for each origin type in each year, $\overline{Earn}_o^{y+1} - \overline{Earn}_o^y$.⁴⁴ We then run the following regression at the o -type level for each year $y \in [1993 - 2011]$:

$$\overline{Earn}_o^{y+1} - \overline{Earn}_o^y = \beta_0^y + \beta_1^y (\ln(C_o^{CF,y+1}) - \ln(C_o^{CF,y})) + \nu_o^y \quad (52)$$

Recall that the $\bar{r}_o^{CF,y}$ values represent predicted money metric utility gains, and are thus denominated in dollars. However, even if the surplus values $\{\theta_g\}$ are time invariant over the chosen period (and the other assumptions of the assignment model specified above all hold, including the approximations just described), dollar-valued mean utility gains would not equal mean annual earnings gains for a given origin type if its workers systematically moved to jobs featuring better or worse amenities (the term $f^i(A(j))$ in the utility function (1) above), avoided more moving/recruiting training costs $c^i(j, k)$, or moved to jobs featuring better or worse continuation values. However, if such changes in other sources of utility nearly cancel out among workers assigned to the same origin type (for all origin types), then $\bar{r}_o^{CF,y}$ should approximately equal $\overline{Earn}_o^{y+1} - \overline{Earn}_o^y$. This implies that $\beta_1^y \approx \sigma^y$.

Clearly, given the additional strong assumptions required, this approach represents a relatively crude attempt to calibrate σ . Indeed, further efforts could conceivably be taken to exclude origin types o' whose surplus values $\{\theta_g : o(g) = o'\}$ were known to be changing over the chosen time period, or to allow θ_g to evolve in a particular parametric fashion.⁴⁵ In fact, Galichon and Salanié (2015) discuss how a vector of σ values associated with different types or combinations of types based on observed characteristics might potentially be

⁴⁴Note that while worker earnings in origin job matches were used to assign workers to skill categories, to this point we have not used observed worker earnings in destination positions to identify any other parameters.

⁴⁵In our actual implementation, we do allow the set of θ_g used to generate the counterfactual prediction to evolve over time in an extremely restricted fashion: we allow the relative payoff of retaining existing workers relative to hiring new workers to evolve over time to match the share of workers who stay at their dominant jobs in each observed year. We do this because the well-chronicled decline in job-to-job mobility during this time period is strongly at odds with the assumption that θ_g is completely time invariant.

jointly estimated with other model parameters (thereby allowing heteroskedasticity across types in the idiosyncratic match component). Since our focus is primarily on examining *relative* incidence across different origin types from shocks featuring different changes in labor demand composition, we opted for the simpler, more transparent approach. In practice, the estimates we obtain for σ^y are fairly consistent across years. We used the mean estimate across all years, $\bar{\sigma} = 8,430$, to produce dollar values for all the results relating to utility gains presented in the paper.

A4 Using Transfers to Decompose the Joint Surpluses $\{\theta_g\}$

This appendix examines whether observing equilibrium transfers, denoted w_{ik} , allows the identification of additional parameters of interest. In Choo and Siow (2006)'s assignment model, the unobserved match-level heterogeneity is assumed to take the form $\epsilon_{ijk} = \epsilon_{o(i,j)k}^1 + \epsilon_{ijd(k)}^2$, so that aggregate surplus is left unchanged when two pairs of job matches (i, k) and (i', k') belonging to the same group g swap partners. The elimination of any true (i, k) match-level surplus component implies that equilibrium transfers cannot vary among job matches belong to the same group g , so that $w_{ik} = w_{g(i,k)} \forall (i, k)$.⁴⁶ Galichon and Salanié (2015) show that under this assumption about the form of unobserved heterogeneity, observing the (common) group-level transfers w_g would be sufficient to decompose the group-level mean joint surplus θ_g into the worker and position's respective pre-transfer payoffs, which we denoted θ_g^l and θ_g^f , respectively.

Because the model proposed in section 2.3 does not impose the additive separability assumption $\epsilon_{ijk} = \epsilon_{o(i,j)k}^1 + \epsilon_{ijd(k)}^2$, equilibrium transfers will in general vary among (i, k) pairs belonging to the same group g . Indeed, given that we observe substantial earnings variance within observed groups g regardless of the worker, firm, and job transition characteristics used to define g , the Choo and Siow (2006) restriction on the nature of unobserved match-level heterogeneity would be strongly rejected in the labor market context.

However, we can still consider the value for identification of the observed transfers $\{w_{ik}\}$. Recall from section 2.2 that equilibrium transfers are related to equilibrium worker and firm payoffs via:

$$w_{ik} = \pi_{ik}^f - q_k \tag{53}$$

$$w_{ik} = r_i - \pi_{ik}^l \tag{54}$$

Next, recall from equation (22) that under Assumptions 1 and 2 the log odds that a randomly chosen position from arbitrary destination type d will choose a worker whose hire

⁴⁶If $w_{ik} > w_{i'k'}$ for any two matched pairs (i, k) and (i', k') such that $g(i, k) = g(i', k')$, then (i', k) would form a blocking pair by proposing a surplus split between them featuring a transfer between w_{ik} and $w_{i'k'}$, thus undermining the stability of the proposed matching.

would be assigned to group g_1 relative to g_2 are given by:

$$\ln\left(\frac{P(g_1|d)}{P(g_2|d)}\right) = \ln(P(g_1|d)) - \ln(P(g_2|d)) = \frac{\theta_{g_1}}{\sigma} + \ln(P(z(g_1)|o(g_1), d)) + \ln(f(o(g_1))) + \ln(C_{o(g_1)}) - \frac{\theta_{g_2}}{\sigma} - \ln(P(z(g_2)|o(g_2), d)) - \ln(f(o(g_2))) - \ln(C_{o(g_2)}) \quad (55)$$

Since $\ln(P(z(g_1)|o(g_1), d))$, $\ln(P(z(g_2)|o(g_2), d))$, $\ln(f(o(g_1)))$, and $\ln(f(o(g_2)))$ are all observed (or, if a large sample is taken, extremely precisely estimated), we can instead form adjusted log odds as in (23):

$$\ln\left(\frac{\hat{P}_{g_1|d}/(P(z(g_1)|o(g_1), d)f(o(g_1)))}{\hat{P}_{g_2|d}/(P(z(g_2)|o(g_2), d)f(o(g_2)))}\right) = \left(\frac{\theta_{g_1} - \theta_{g_2}}{\sigma}\right) + (\ln(C_{o(g_1)}) - \ln(C_{o(g_2)})) \quad (56)$$

Under Assumption 1, C_o is the mean of exponentiated (and rescaled) equilibrium utility payoffs owed to workers $i : o(i) = o$:

$$C_o = \frac{1}{|o|} \sum_{i:o(i,j(i))=o} e^{-\frac{r_i}{\sigma}} \approx \sum_{\frac{1}{g_k}} \sum_{i:g(i,j(i),k)=g} e^{-\frac{r_i}{\sigma}} \quad \forall k \quad (57)$$

Plugging (54) into (57) and then (57) into (56) yields:

$$\begin{aligned} & \ln\left(\frac{\hat{P}_{g_1|d}/(P(z(g_1)|o(g_1), d)f(o(g_1)))}{\hat{P}_{g_2|d}/(P(z(g_2)|o(g_2), d)f(o(g_2)))}\right) \\ &= \left(\frac{\theta_{g_1} - \theta_{g_2}}{\sigma}\right) + \left(\ln\left(\frac{1}{|o|} \sum_{i:o(i,j(i))=o(g_1)} e^{-\frac{w_{ik} + \pi_{ik}^l}{\sigma}}\right) - \ln\left(\frac{1}{|o|} \sum_{i:o(i,j(i))=o(g_2)} e^{-\frac{w_{ik} + \pi_{ik}^l}{\sigma}}\right)\right) \quad (58) \end{aligned}$$

It is not immediately obvious how to use equation (60) to recover parameters of interest. Only when we add further assumptions that are at odds with the structure of the model can we recover an expression that mirrors the one in Choo and Siow (2006). Specifically, suppose the following assumptions hold:

$$\begin{aligned} r_i &\approx r_{o(i)} \quad \forall i : o(i, j(i)) = o \quad \forall o \in \mathcal{O} \\ \pi_{ik}^l &= \pi_{g(i,k)}^l \equiv \theta_g^l \quad \forall (i, k) : g(i, k) = g \quad \forall g \in \mathcal{G} \\ w_{ik} &= w_{g(i,k)} \quad \forall (i, k) : g(i, k) = g \quad \forall g \in \mathcal{G} \quad (59) \end{aligned}$$

We suspect that these assumptions will be extremely unlikely to hold in any stable matching if there is meaningful variance in the unobserved match surplus component $\epsilon_{ij(i)k}$ among (i, k) pairs assigned to the same group g , whose extreme value distribution is the basis for the logit closed-forms used for conditional choice probabilities above. Nonetheless, they

yield:

$$\begin{aligned}
& \ln\left(\frac{\hat{P}_{g_1|d}/(P(z(g_1)|o(g_1), d))f(o(g_1))}{\hat{P}_{g_2|d}/(P(z(g_2)|o(g_2), d))f(o(g_2))}\right) \\
&= \left(\frac{\theta_{g_1} - \theta_{g_2}}{\sigma}\right) + (\ln(e^{-r_{o(g_1)}}) - \ln(e^{-r_{o(g_2)}})) \\
&= \left(\frac{\theta_{g_1} - \theta_{g_2}}{\sigma}\right) + \frac{-r_{o(g_1)} + r_{o(g_2)}}{\sigma} \\
&= \left(\frac{\theta_{g_1} - \theta_{g_2}}{\sigma}\right) + \left(\frac{-(w_{g_1} + \theta_{g_1}^l) + (w_{g_2} + \theta_{g_2}^l)}{\sigma}\right) \\
&= \frac{\theta_{g_1}^f - \theta_{g_2}^f + (w_{g_2} - w_{g_1})}{\sigma} \tag{60}
\end{aligned}$$

Given an estimate of σ based on multiple markets (as described in Appendix A3) and data on mean annual earnings for each transition group $g \in \mathcal{G}$, one could identify the difference in the position component of the joint surplus for arbitrary groups g_1 and g_2 . This provides information about the relative profit contributions of different types of workers for each type of firm before such workers salaries are considered. Note that one could still not separate the training cost, recruiting cost, current revenue contribution, and continuation value components of θ_g^f without additional data.

A similar progression using adjusted log odds based on the worker side conditional probabilities $P(g_1|o_1)$ and $P(g_2|o_1)$ would yield an estimate of the corresponding difference in the worker components of the joint surplus $\theta_{g_1}^l - \theta_{g_2}^l$ for any two groups featuring the same origin worker type. Since one such group could represent nonemployment, this approach would provide estimates of the desirability of working at various types of firms in various locations for zero pay relative to nonemployment. These values identify the reservation salary necessary to convince each origin worker type to take (or continue) a position of each destination type. As with firms, one could not disentangle the moving cost, search cost, non-wage amenity value, and continuation value components of the surplus without further data.

Because 1) we deem the assumptions (59) to be antithetical to the spirit of the model and at odds with the data, and 2) other than estimating σ , the use of transfers is not necessary to fulfill the primary aim of the paper, evaluating the utility and profit incidence across worker and position types of alternative local labor demand shocks, we do not make further use of the observed annual earnings distributions in the destination period y in any aggregate labor market transition $(y - 1, y)$ in this paper.

A5 Smoothing Procedure

In this appendix we describe how we smooth the empirical distribution of transitions across transition groups, $\hat{P}(g)$, prior to estimation in order to generate accurate estimates of the

elements of the identified set of joint surplus difference-in-differences Θ^{D-in-D} . We smooth for two reasons. First, such smoothing serves as a “noise infusion” technique that removes the risk that the identity of any particular individual or firm could be revealed by any of the estimates presented in the paper, as required of any research results generated from confidential microdata in Federal Statistical Research Data Centers (FSRDCs). Second, smoothing is necessary because there are sufficiently few observations per transition group that many transition groups are either rarely (or never) observed in a given cross-sectional transition despite substantial underlying matching surpluses simply due to sampling error. Essentially, $\hat{P}(g)$ is only a consistent estimator of $P(g)$ as the number of observed worker transitions per group I/G approaches infinity.

We overcome this sampling error problem by assuming that the underlying frequency $P(g)$ with which a randomly chosen transition belongs to a particular transition group is a smooth function of the observed characteristics that define group g . This permits the use of a kernel density estimator that computes a weighted average of the empirical probabilities $\hat{P}(g')$ of “nearby” groups g' that feature “similar” vectors of characteristics to generate a well-behaved approximation of $P(g)$ from the noisy empirical distribution $\hat{P}(g)$.

Such smoothing introduces two additional challenges. First, excessive smoothing across other transition groups erodes the signal contained in the data about the degree of heterogeneity in the relative surplus from job transitions featuring different combinations of worker characteristics, firm characteristics, and origin and destination locations. Since highlighting the role of such heterogeneity in forecasting the incidence of labor market shocks is a primary goal of the paper, decisions about the appropriate bandwidth must be made with considerable thought. The second, related challenge consists of identifying which of the worker and firm characteristics that defines other groups makes them “similar”, in the sense that the surplus $\{\theta_{g'}\}$ is likely to closely approximate the surplus whose estimate we wish to make more precise θ_g .

Recall that each group $g \equiv g(o, d, z)$ is a combination of 1) the origin establishment location (which we denote $loc(o)$) and workers’ initial earnings quintile (or nonemployment status) at the origin establishment (denoted $earn(o)$); 2) the destination establishment’s location ($loc(d)$), firm size category ($f_size(d)$), firm average earnings category ($f_earn(d)$), and industry supersector ($ind(d)$); and 3) the indicator $z(i, j, k)$ for whether establishment j and establishment k are the same, so that worker i is a job stayer rather than a mover (denoted $stayer(g)$).

Given our goal of accurately characterizing incidence at a very low level of geographic aggregation, we wish to preserve as accurately as possible any signal in the data about the structure of spatial ties between nearby local areas. Thus, wherever possible our kernel estimator should place non-zero weight only on alternative groups g' that share the same origin and destination locations ($loc(o(g)) = loc(o(g'))$ and $loc(d(g)) = loc(d(g'))$). Similarly, we suspect that the combination of the non-location characteristics firm size, firm

average worker earnings, and firm industry is likely to be more important than location in determining the skill category of worker (proxied by initial earnings quintile) that generates the most surplus. To develop a smoothing approach that embodies these principles, we exploit the fact that $P(g)$ can be decomposed via:

$$\begin{aligned}
P(g) &= P(g|d(g))h(d(g)) = P([o(g), d(g), z(g)]|d)h(d(g)) \\
&= P([loc(o(g)), earn(o(g)), stayer(g)]|d)h(d(g)) \\
&= P(loc(o(g))|earn(o(g)), stayer(g), d)P([earn(o(g)), stayer(g)]|d)h(d(g)) \\
&= 1(stayer(g) = 1)P(loc(o(g))|earn(o(g)), 1(stayer(g) = 1), d)P([earn(o(g)), 1(stayer(g) = 1)]|d)h(d(g)) \\
&+ 1(stayer(g) = 0)P(loc(o(g))|earn(o(g)), 1(stayer(g) = 0), d)P([earn(o(g)), 1(stayer(g) = 0)]|d)h(d(g)) \\
&= 1(stayer(g) = 1)1(loc(o(g)) = loc(d(g)))P([earn(o(g)), 1(stayer(g) = 1)]|d)h(d(g)) \\
&+ 1(stayer(g) = 0)P(loc(o(g))|earn(o(g)), 1(stayer(g) = 0), d)P([earn(o(g)), 1(stayer(g) = 0)]|d)h(d(g))
\end{aligned} \tag{61}$$

where the first two lines use the law of total probability and the set of characteristics that define $o(g)$ and $z(g)$, the third line uses the fact that the $z(g) \equiv Stayer(g)$ only takes on two values (0 for job movers and 1 for job stayers), and the last line uses the fact that $P(loc(o(g))|earn(o(g)), 1(stayer(g) = 1), d) = 1(loc(o(g)) = loc(d(g)))$, since a potential stayer associated with a particular destination type must have already been working at the same location in the origin period (by virtue of being a job stayer, since we treat firms that switch locations as different firms for computational reasons).

We use separate kernel density estimator procedures to estimate each of $P(loc(o(g))|earn(o(g)), 1(stayer(g) = 0), d(g))$, $P(earn(o(g)), 1(stayer(g) = 0)|d(g))$, and $P(earn(o(g)), 1(stayer(g) = 1)|d(g))$.

Consider first the estimation of $P(loc(o(g))|earn(o(g)), 1(stayer(g) = 0), d(g))$, the conditional probability that a particular new hire would be originally located at location $loc(o)$, given the hired worker's initial earnings category and the destination type d of the hiring position. Let $K^{dist}(g, g')$ represent the metric capturing how similar an alternative group g' is to g for the purpose of estimating the propensity for firms of type d to hire workers from a particular location (conditional on skill level). As discussed above, wherever possible we only assign non-infinite distance $K^{dist}(g, g') < \infty$ (which corresponds to non-zero weight) to empirical conditional probabilities $P(loc(o(g'))|earn(o(g')), 1(stayer(g') = 0), d(g'))$ of alternative groups g' that feature both the same origin location $loc(o(g')) = loc(o(g))$ and the same destination location $loc(d(g')) = loc(d(g))$.⁴⁷

$K^{dist}(g, g')$ assigns the smallest distance to alternative groups g' that also feature the same destination type ($d(g') = d(g)$), so that g and g' only differ in the initial earnings

⁴⁷There are a very small number of destination and origin types are never observed in any transition. By necessity, we put positive weight on groups featuring nearby origin or destination locations in such cases.

category of hired workers. The closer $earn(o(g'))$ is to $earn(o(g))$, the smaller is the assigned distance $K^{dist}(g, g')$, but the profile flattens so that all groups g' that differ from g only due to $earn(o(g'))$ contribute to the weighted average. $K^{dist}(g, g')$ assigns larger (but still noninfinite) distance to groups g' featuring destination types that also differ on firm size, firm avg. earnings, or industry dimensions. The more different the firm composition of the group, the smaller is its weight, with the profile again flattening so that all groups g' featuring the same origin and destination locations receive non-zero weight. Thus, groups with less similar worker and firm characteristics receive non-negligible weight only when there are too few observations from groups featuring more similar worker and firm characteristics to form reliable estimates. The weight assigned to a particular alternative group g' also depends on the number of observed new hires made by $d(g')$ at a particular skill level $earn(o(g'))$, denoted $N^{dist}(g')$ below, since this determines the signal strength of the empirical conditional choice probability $P(loc(o(g'))|earn(o(g')), 1(stayer(g') = 0), d(g'))$. Thus, we have:

$$P(loc(o(g))|earn(o(g)), 1(stayer(g) = 0), d(g)) \approx \sum_{g'} \left(\frac{\phi(K^{dist}(g', g)N^{dist}(g'))}{\sum_{g''} \phi(K^{dist}(g'', g)N^{dist}(g''))} \hat{P}(loc(o(g'))|earn(o(g')), 1(stayer(g') = 0), d(g')) \right) \quad (62)$$

where $\phi(*)$ is the normal density function (used as the kernel density), and $\frac{\phi(K^{dist}(g', g)N^{dist}(g'))}{\sum_{g''} \phi(K^{dist}(g'', g)N^{dist}(g''))}$ represents the weight given to a particular nearby transition group g' .

Next, consider the estimation of $P(earn(o(g)), 1(stayer(g) = 1)|d)$ and $P(earn(o(g)), 1(stayer(g) = 0)|d)$, the conditional probabilities that either a job stayer or mover originally paid at a particular earnings quintile (or possibly non-employed for movers) will be hired to fill a position of destination type d . Let $K^{earn/move}(g, g')$ and $K^{earn/stay}(g, g')$ represent the metrics capturing how similar alternative groups g' are to g for the purpose of estimating the propensity for firms of type d to hire (or retain) workers at particular skill levels.

$K^{earn/move}(g, g')$ and $K^{earn/stay}(g, g')$ each assign infinite distance (translating to zero weight) to groups g' featuring different combinations of firm size, average worker earnings, or industry than the target group g . $K^{earn/move}(g, g')$ ($K^{earn/stay}(g, g')$) assigns small distances to the conditional probabilities associated with groups g' representing hiring new (retaining) workers from the same initial earnings (or nonemployment) category $earn(o(g)) = earn(o(g'))$ among firms from the same destination type $d(g) = d(g')$ but who are hiring workers from nearby locations. The distance metric increases in the tract pathlength between $loc(o(g'))$ and $loc(o(g))$, but flattens beyond a threshold distance, so that groups featuring all origin worker locations (but same other characteristics) contribute to the estimate.

Larger (but finite) distance values for $K^{earn/move}(g, g')$ and $K^{earn/stay}(g, g')$ are assigned to conditional probabilities from groups g' that feature different (but nearby) des-

tionation locations (so $d(g) \neq d(g')$) but the same combination of firm size quartile, firm average worker earnings quartile, and industry supersector. Again, the distance metric increases in the pathlength between $loc(d(g))$ and $loc(d(g'))$, but eventually flattens at a large but non-infinite value. As before, the weight given to a group g' also depends on the precision of its corresponding number of total hires made by firms of the destination type $d(g')$, which is proportional to $h(d(g'))$.

Again, the motivation here is that targeted skill level and the decision to retain workers vs. hire new workers (conditional on the utilities bids required by workers in different locations) is likely to be driven to a greater extent by the type of production process (as proxied by size, mean worker earnings, and industry) than by the location of the firm. Nonetheless, since we still suspect that there is unobserved heterogeneity in production processes conditional on our other firm observables that might be spatially correlated, we place greater weight on the skill/retention decisions of geographically proximate firms. More distant firms receive non-negligible weight only when there are too few local observations to form reliable estimates. The estimators for $P(earn(o(g)), 1(stayer(g) == 1)|d)$ and $P(earn(o(g)), 1(stayer(g) == 0)|d)$ can thus be represented via:

$$P(earn(o(g)), 1(stayer(g) == 0)|d(g)) \approx \sum_{g'} \left(\frac{\phi(K^{earn/move}(g', g)h(d(g')))}{\sum_{g''} \phi(K^{earn/move}(g'', g)h(d(g'')))} \hat{P}(earn(o(g')), 1(stayer(g') = 0)|d(g')) \quad (63)$$

$$P(earn(o(g)), 1(stayer(g) = 1)|d(g)) \approx \sum_{g'} \left(\frac{\phi(K^{earn/stay}(g', g)h(d(g')))}{\sum_{g''} \phi(K^{earn/stay}(g'', g)h(d(g'')))} \hat{P}(earn(o(g')), 1(stayer(g') = 1)|d(g')) \quad (64)$$

Bringing the pieces together, this customized smoothing procedure has a number of desirable properties. First, by requiring the same origin and destination locations as a necessary condition for non-zero weight when estimating the propensity for particular destination types to hire workers from each location, we can generate considerable precision in estimated conditional choice probabilities without imposing any assumption about the spatial links between locations. Second, at the same, we can still use information contained in the hiring and retention choices of more distant firms to learn about the propensity for firms of different size, pay level, and industry to retain and hire workers at different skill levels and from nonemployment. Third, the procedure places non-trivial weight on transition groups featuring less similar worker and firm characteristics only when there are too few observed hires/retentions made by firms associated with groups featuring very similar characteristics to yield reliable estimates. Fourth, overall the estimated probabilities $P(g|d)$ place weight on many different groups, so that no element of the resulting smoothed group-level distribution contains identifying information about any particular worker or firm, eliminating disclosure risk.

A6 Imputing Missing NE-to-NE, E-to-NE, and NE-to-E Transitions

This appendix discusses the procedure used to impute missing nonemployment-to-nonemployment (hereafter NE-to-NE) transitions, employment-to-nonemployment (hereafter E-to-NE) transitions, and nonemployment-to-employment (hereafter NE-to-E) transitions caused by the combination of our decisions to drop new market entrants and retirees and the existence of incomplete employment histories due to a limited number of available years of data.

As discussed in the main text, we drop from the sample the first NE-to-E transition in each worker's history on the basis that initial market entrants are likely to generate very different surplus from their transitions than other non-employed workers, and we have no way of determining how long the initial nonemployment spell that preceded their first jobs lasted.⁴⁸ Thus, if such transitions are left in the sample, we may overstate the propensity for firms to hire nonemployed workers, thereby biasing the potential impact of local labor demand shocks on nonemployed workers. However, in the early years of the sample, the first NE-to-E transition observed may not be the worker's initial labor market entry; rather, their previous employment may have occurred prior to the first sample year. Indeed, in the first year of the sample, any worker who was nonemployed in the first year will be treated as a labor market entrant the second year. This problem dissipates as the sample progresses, so that by the end the only improperly dropped NE-to-E transitions will be those whose nonemployment spell between years with primary jobs lasted nearly twenty years. Thus, we can use the prevalence of NE-to-E transitions in which the worker is previously observed working in the later years of the sample to impute the number of inappropriately dropped NE-to-E transitions in the early sample years in which the worker's prior year of employment is unobserved. We then account for true business cycle variation in the frequency of NE-to-E transitions (as well as state-year combinations in which data are incomplete) by incorporating the evolution of the frequency of employment-to-employment (hereafter E-to-E) transitions (which are always included in the sample) across state-year combinations.

In the same vein, we also drop from the sample the last E-to-NE transition in each worker's history, again on the basis that retirees are likely to generate different surplus from their transitions than other newly nonemployed workers, and we have no way of determining how many "retirees" actually were unemployed for multiple years prior to their last job before giving up their job searches. The concern is that keeping retirees in the sample artificially increases the frequency with which relatively high paid workers transi-

⁴⁸An alternative approach would be to simply define workers as potential employees during a particular age window, say 18-70. This would remove the problem of not observing the unemployment spell associated with the initial job search, but in the absence of data on worker education would still lump those with freshly minted college degrees with less educated workers who had been unemployed the previous year, two groups with very different job finding rates.

tion to non-employment, which may lead to an overstatement of the degree to which the nonemployment incidence of a natural disaster falls on more skilled workers (for example). However, as with labor market entrants, by dropping each worker’s final *observed* E-to-NE transition, we risk labeling as retirees workers who are merely facing a temporary spell of nonemployment, and who will return to employment in the years following the end of the sample. The severity of this problem is greatest at the end of the sample, with all final year nonemployed workers being treated as retirees. Our solution to this problem is analogous to the one for labor market entrants: use the prevalence of E-to-NE transitions in which the worker is later observed working in the early years of the sample to impute the number of inappropriately dropped E-to-NE transitions in the later sample years in which the workers’ future years of employment are unobserved (augmented again with time series variation in E-to-E transitions).

Finally, NE-to-NE transitions are only kept in the sample when a worker is observed as employed at a qualifying primary job in both a prior and future year. Thus, truncated employment histories at both the front and back ends of the sample contribute to understatement of the frequency of NE-to-NE transitions. The prevalence of NE-to-NE transitions in the middle of the sample is used to impute the number of NE-to-NE transitions that were inappropriately dropped due to truncated employment histories.

Consider NE-to-E transitions first. We impute the number of inappropriately dropped NE-to-E transitions associated with each relevant transition group g featuring nonemployment as part of the origin type and employment at some job as part of the destination type as follows:

1. Compute the number of NE-to-E and E-to-E transitions remaining in the sample for each state-tract-age group-year combination.
2. Construct the ratio of (NE-to-E)/(E-to-E) transitions for each state-age-year combination, denoted $Ratio_{say}$.
3. For each state-age group combination, identify the maximum value of $Ratio_{say}$ among the sample years: $Ratio_{sa}^{peak} = \max_{y' \in [1990, 2010]} Ratio_{say'}$. We treat this year’s ratio as the “true” ratio for the identified year for each state-age group combination.
4. Divide $Ratio_{say}/Ratio_{sa}^{peak}$ to convert the NE-to-E/E-to-E ratio to a share of the peak ratio for the given state-age group, denoted $Share_{say}$. This effectively removes the persistent variation in $Ratio_{say}$ across state-age group combinations.
5. Regress $Share_{say}$ on a set of dummies indicating the combination of age group and number of years between y and the last year of the sample (2010), $1(2010 - y = x, agegroup = z)$, for $(x, z) \in [0, 20] \times [1, 10]$, and a set of dummies indicating the combination of age group and number of years between y and the first year state

s enters the sample, $1(y - Firstyear(s) = x, agegroup = z, \text{ for } (x, z) \in [0, 2010 - Firstyear(s)] \times [1, 10]$. Because each state exits the sample in the same year, the first set of dummies are equivalent to age group \times year dummies. Because different states enter the sample in different years between 1990 and 1993, the second set of dummies are not collinear with the first set. Denote the predicted values by \hat{Share}_{say} . The predicted values from this regression capture the propensity for NE-to-E transitions to be undercounted relative to E-to-E transitions in the early years of the sample. By only using the number of years until the end of the sample and the timing of states entry into the sample to generate the prediction, we seek to isolate the component of variation in $Share_{say}$ that is attributable to the selection bias generated by the way we define the sample, thereby removing the part of the variation in $Share_{say}$ that is due to differences in the timing and severity of economic booms and busts across states that generate true variation in $Share_{say}$ unrelated to our sample selection procedure. Interacting with age groups allows the severity of the bias to be different for different age groups due to a decreased likelihood of unobserved prior employment spells for younger workers.

6. Construct the expected number of “missing” NE-to-E transitions for each state-age-year combination as:

$$NEtoE_{say}^{imputed} = (EtoE_{say})Ratio_{sa}^{peak}(1 - \hat{Share}_{say}) \quad (65)$$

$(1 - \hat{Share}_{say})$ captures the NE-to-NE transitions that are missing due to our sample selection procedure as a share of E-to-E transitions, relative to the natural share for the chosen state-age group combination (as judged by $Ratio_{sa}^{peak}$). Multiplying this value by $Ratio_{sa}^{peak}$ converts this relative measure back to a ratio that is specific to the particular state-age group context, and multiplying by $EtoE_{say}$ converts this ratio into a count of missing NE-to-E transitions.

7. Finally, we distribute the imputed NE-to-E transitions for each state-age-year combination across transition groups g featuring different census tracts within the chosen state and different destination types by multiplying $NEtoE_{say}^{imputed}$ by the average share of all observed NE-to-E transitions in the chosen state-age group combination (across all years) associated with each transition group, denoted $P(g|s(g), a(g))$:

$$NEtoE_g^{imputed} = (NEtoE_{say(g)}^{imputed})(P(g|s(g), a(g))) \quad (66)$$

Analogous imputation procedures are used for E-to-NE and NE-to-NE transitions. For E-to-NE transitions, most imputed transitions are at the end of the sample, and the state-age-year imputed counts $EtoNE_{say}^{imputed}$ need to be distributed across origin locations and initial earnings levels only, since non-employment in year y is its own destination type

d. Similarly, the bulk of the imputed NE-to-NE transitions are at the beginning and end of the sample, with few in the middle years, and the state-age-year imputed counts $NEtoNE_{say}^{imputed}$ only need to be distributed across origin locations (since for observed NE-to-NE transitions we assign the origin location to be the geographic location of the most recent establishment at which a worker worked).

Tables

Table A1: Probability of Obtaining Stimulus Job for a Randomly Chosen Individual at Different Distances from Focal Tract: Stimuli Consist of 500 New Jobs in Different Industries (Averaged Across Firm Size/Firm Average Earnings Combinations)

Distance from Focal Tract	Industry								
	Avg.	Info.	Manu.	R/W Trd.	Oth. Serv.	Ed./Hlth	Lei/Hosp.	Gov.	Const.
Target Tract	0.032	0.028	0.034	0.031	0.036	0.030	0.031	0.028	0.035
1 Tct Away	0.005	0.004	0.005	0.004	0.006	0.005	0.005	0.005	0.005
2 Tcts Away	0.002	0.002	0.002	0.002	0.003	0.003	0.002	0.002	0.002
3+ Tcts w/in PUMA	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001
1 PUMA Away	5.7E-04	5.4E-04	5.6E-04	5.4E-04	6.2E-04	6.6E-04	5.9E-04	5.3E-04	5.5E-04
2 PUMAs Away	2.5E-04	2.5E-04	2.5E-04	2.6E-04	2.7E-04	2.7E-04	2.6E-04	2.5E-04	2.4E-04
3+ PUMAs w/in State	1.3E-04	1.3E-04	1.4E-04	1.5E-04	1.2E-04	1.3E-04	1.3E-04	1.3E-04	1.2E-04
1-2 States Away	2.7E-05	2.6E-05	2.8E-05	3.1E-05	2.8E-05	2.5E-05	2.8E-05	2.4E-05	2.4E-05
3+ States Away	1.3E-06	9.6E-07	1.8E-06	1.5E-06	9.1E-07	1.6E-06	1.6E-06	1.0E-06	9.4E-07

Notes:

Table A2: Share of Stimulus Positions Filled by Workers Initially Employed (or Nonemployed) at Different Distances from Focal Tract:
 Stimuli Consist of 500 New Jobs in Different Industries (Averaged Across Firm Size/Firm Average Earnings Combinations)

Distance from Focal Tract	Industry								
	Avg.	Info.	Manu.	R/W Trd.	Oth. Serv.	Ed./Hlth	Lei/Hosp.	Gov.	Const.
Target Tract	0.075	0.069	0.077	0.070	0.083	0.067	0.070	0.071	0.090
1 Tct Away	0.067	0.063	0.066	0.056	0.076	0.070	0.066	0.071	0.067
2 Tcts Away	0.073	0.076	0.070	0.063	0.079	0.074	0.071	0.074	0.074
3+ Tcts w/in PUMA	0.136	0.131	0.132	0.125	0.136	0.144	0.141	0.138	0.143
1 PUMA Away	0.107	0.108	0.102	0.098	0.112	0.119	0.106	0.106	0.109
2 PUMAs Away	0.131	0.136	0.123	0.127	0.131	0.132	0.128	0.138	0.129
3+ PUMAs w/in State	0.101	0.104	0.101	0.112	0.090	0.097	0.098	0.107	0.102
1-2 States Away	0.249	0.261	0.249	0.280	0.252	0.224	0.248	0.240	0.235
3+ States Away	0.061	0.051	0.080	0.070	0.041	0.073	0.072	0.055	0.049

Notes:

Table A3: Change in Probability of Nonemployment due to Stimulus for a Randomly Chosen Individual at Different Distances from Focal Tract: Stimuli Consist of 500 New Jobs at Firms in Alternative Industries (Averaged Across Firm Size/Firm Average Earnings Combinations)

Distance from Focal Tract	Industry								
	Avg.	Info.	Manu.	R/W Trd.	Oth. Serv.	Ed./Hlth	Lei/Hosp.	Gov.	Const.
Target Tract	0.012	0.009	0.011	0.011	0.014	0.011	0.011	0.010	0.014
1 Tct Away	0.002	0.001	0.001	0.001	0.002	0.002	0.002	0.002	0.002
2 Tcts Away	8.7E-04	8.2E-04	8.3E-04	8.1E-04	1.1E-03	9.7E-04	8.8E-04	8.3E-04	7.9E-04
3+ Tcts w/in PUMA	5.2E-04	4.8E-04	4.9E-04	4.9E-04	5.7E-04	5.9E-04	5.5E-04	5.0E-04	5.2E-04
1 PUMA Away	3.6E-04	3.3E-04	3.5E-04	3.5E-04	4.0E-04	3.9E-04	3.7E-04	3.5E-04	3.5E-04
2 PUMAs Away	2.2E-04	2.1E-04	2.1E-04	2.2E-04	2.3E-04	2.3E-04	2.2E-04	2.1E-04	2.1E-04
3+ PUMAs w/in State	1.5E-04	1.4E-04	1.5E-04	1.5E-04	1.5E-04	1.5E-04	1.5E-04	1.4E-04	1.4E-04
1-2 States Away	4.2E-05	4.1E-05	4.3E-05	4.5E-05	4.4E-05	4.2E-05	4.3E-05	4.0E-05	3.9E-05
3+ States Away	3.7E-06	3.3E-06	4.5E-06	4.1E-06	3.6E-06	4.0E-06	4.0E-06	3.2E-06	3.3E-06

Notes:

Table A4: Share of Additional Employment Produced by Stimulus Among Geographic Areas Defined by Distances from the Focal Tract:
 Stimuli Consist of 500 New Jobs at Firms in Different Industries (Averaged Across Firm Size/Firm Average Earnings Combinations)

Distance from Focal Tract	Industry								
	Avg.	Info.	Manu.	R/W Trd.	Oth. Serv.	Ed./Hlth	Lei/Hosp.	Gov.	Const.
Target Tract	0.027	0.023	0.025	0.025	0.031	0.024	0.025	0.025	0.035
1 Tct Away	0.021	0.020	0.019	0.017	0.025	0.024	0.022	0.022	0.022
2 Tcts Away	0.026	0.026	0.024	0.023	0.030	0.028	0.025	0.026	0.025
3+ Tcts w/in PUMA	0.067	0.065	0.061	0.061	0.071	0.073	0.069	0.068	0.072
1 PUMA Away	0.068	0.066	0.063	0.063	0.072	0.071	0.068	0.070	0.069
2 PUMAs Away	0.111	0.113	0.105	0.109	0.114	0.113	0.111	0.115	0.112
3+ PUMAs w/in State	0.114	0.117	0.111	0.114	0.111	0.113	0.113	0.119	0.116
1-2 States Away	0.392	0.406	0.389	0.405	0.389	0.378	0.388	0.395	0.384
3+ States Away	0.177	0.167	0.202	0.185	0.161	0.182	0.183	0.164	0.171

Notes:

Table A5: Change in Probability of Nonemployment due to Stimulus for a Randomly Chosen Individual at Different Distances from Focal Tract: Stimuli Consist of 500 New Positions in Alternative Combinations of Firm Size Quartile/Firm Average Pay Quartile (Averaged Across Industry Supersectors)

Distance from Focal Tract	Firm Size/Pay Level Combination			
	Sm./Low	Lg./Low	Sm./Hi	Lg./Hi
Target Tract	0.013	0.013	0.009	0.011
1 Tct Away	0.002	0.002	0.001	0.002
2 Tcts Away	9.3E-04	1.0E-03	6.9E-04	8.7E-04
3+ Tcts w/in PUMA	5.5E-04	5.9E-04	4.2E-04	5.3E-04
1 PUMA Away	3.8E-04	4.0E-04	2.9E-04	3.8E-04
2 PUMAs Away	2.2E-04	2.4E-04	1.8E-04	2.3E-04
3+ PUMAs w/in State	1.5E-04	1.6E-04	1.2E-04	1.6E-04
1-2 States Away	4.3E-05	4.3E-05	3.8E-05	4.5E-05
3+ States Away	4.1E-06	3.5E-06	3.7E-06	3.7E-06

Notes:

Table A6: Share of Additional Employment Produced by Stimulus Among Geographic Areas Defined by Distances from the Focal Tract:

Stimuli Consist of 500 New Positions in Alternative Combinations of Firm Size Quartile/Firm Average Pay Quartile (Averaged Across Industry Supersectors)

Distance from Focal Tract	Firm Size/Pay Level Combination			
	Sm./Low	Lg./Low	Sm./Hi	Lg./Hi
Target Tract	0.030	0.029	0.023	0.024
1 Tct Away	0.023	0.024	0.019	0.020
2 Tcts Away	0.027	0.029	0.023	0.025
3+ Tcts w/in PUMA	0.069	0.074	0.061	0.066
1 PUMA Away	0.069	0.072	0.062	0.068
2 PUMAs Away	0.110	0.117	0.105	0.114
3+ PUMAs w/in State	0.112	0.119	0.109	0.117
1-2 States Away	0.383	0.387	0.397	0.401
3+ States Away	0.186	0.158	0.198	0.167

Notes:

Table A7: Expected Change in Utility from New Stimulus Positions for a Randomly Chosen Individual at Different Distances from Focal Tract:
 Stimuli Consist of 500 New Jobs in Different Industries (Averaged Across Firm Size/Firm Average Earnings Combinations)

Distance from Focal Tract	Industry								
	Avg.	Info.	Manu.	R/W Trd.	Oth. Serv.	Ed./Hlth	Lei/Hosp.	Gov.	Const.
Target Tract	512	462	565	493	632	411	552	471	508
1 Tct Away	211	189	197	226	250	207	218	196	204
2 Tcts Away	164	155	156	191	180	152	173	158	147
3+ Tcts w/in PUMA	153	141	145	171	162	149	164	146	141
1 PUMA Away	140	132	128	169	150	135	147	132	128
2 PUMAs Away	115	106	104	148	129	105	125	108	100
3+ PUMAs w/in State	111	101	99	144	123	100	119	103	95
1-2 States Away	92	86	80	125	99	82	102	86	80
3+ States Away	0	0	0	0	0	0	0	0	0

Notes:

Table A8: Expected Share of Utility Increases from New Stimulus Positions by Workers Initially Employed (or Nonemployed) at Different Distances from Focal Tract: Stimuli Consist of 500 New Jobs in Different Industries (Averaged Across Firm Size/Firm Average Earnings Combinations)

Distance from Focal Tract	Industry								
	Avg.	Info.	Manu.	R/W Trd.	Oth. Serv.	Ed./Hlth	Lei/Hosp.	Gov.	Const.
Target Tract	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001
1 Tct Away	0.003	0.003	0.003	0.002	0.003	0.003	0.002	0.003	0.003
2 Tcts Away	0.005	0.005	0.005	0.004	0.005	0.005	0.004	0.005	0.005
3+ Tcts w/in PUMA	0.018	0.018	0.019	0.016	0.018	0.020	0.018	0.019	0.020
1 PUMA Away	0.025	0.025	0.025	0.022	0.024	0.026	0.023	0.025	0.026
2 PUMAs Away	0.054	0.054	0.055	0.052	0.056	0.055	0.053	0.054	0.054
3+ PUMAs w/in State	0.079	0.077	0.079	0.077	0.082	0.079	0.077	0.079	0.078
1-2 States Away	0.784	0.788	0.766	0.799	0.786	0.776	0.793	0.785	0.781
3+ States Away	0.032	0.030	0.047	0.027	0.025	0.035	0.029	0.029	0.032

Notes:

Table A9: Expected Change in Utility From New Stimulus Positions for a Randomly Chosen Individual at Different Distances from Focal Tract:
 Stimuli Consist of 500 New Positions in Alternative Combinations of Firm Size Quartile/Firm Average Pay Quartile (Averaged Across Industry Supersectors)

Distance from Focal Tract	Firm Size/Pay Level Combination			
	Sm./Low	Lg./Low	Sm./Hi	Lg./Hi
Target Tract	546	517	459	525
1 Tct Away	215	230	176	223
2 Tcts Away	163	184	133	176
3+ Tcts w/in PUMA	152	171	124	163
1 PUMA Away	141	158	112	150
2 PUMAs Away	116	135	89	123
3+ PUMAs w/in State	111	129	84	117
1-2 States Away	91	109	70	99
3+ States Away	0	0	0	0

Notes:

Table A10: Share of Additional Employment Produced by Stimulus among Workers Initially Employed (or Nonemployed) at Different Initial Earnings Quintiles (or Nonemployment): Stimuli Consist of 500 New Jobs at Firms in Different Industries (Averaged Across Firm Size/Firm Average Earnings Combinations)

Earnings Quintile	Industry								
	Avg.	Info.	Manu.	R/W Trd.	Oth. Serv.	Ed./Hlth	Lei/Hosp.	Gov.	Const.
Nonemployed	0.433	0.417	0.409	0.429	0.462	0.433	0.430	0.424	0.456
1st Quintile	0.153	0.156	0.153	0.152	0.146	0.158	0.164	0.155	0.143
2nd Quintile	0.147	0.147	0.152	0.147	0.139	0.149	0.149	0.148	0.141
3rd Quintile	0.113	0.115	0.117	0.113	0.108	0.113	0.110	0.116	0.111
4th Quintile	0.086	0.089	0.090	0.087	0.082	0.085	0.083	0.088	0.086
5th Quintile	0.072	0.077	0.077	0.074	0.067	0.069	0.068	0.071	0.070

Notes:

Table A11: Expected Share of Utility Increases from New Stimulus Positions among Workers Initially Employed (or Nonemployed) at Different Initial Earnings Quintiles (or Nonemployment):

Stimuli Consist of 500 New Jobs at Firms in Different Industries (Averaged Across Firm Size/Firm Average Earnings Combinations)

Earnings Quintile	Industry								
	Avg.	Info.	Manu.	R/W Trd.	Oth. Serv.	Ed./Hlth	Lei/Hosp.	Gov.	Const.
Nonemployed	0.138	0.137	0.134	0.140	0.141	0.138	0.139	0.138	0.138
1st Quintile	0.119	0.119	0.118	0.119	0.119	0.119	0.120	0.119	0.118
2nd Quintile	0.166	0.165	0.166	0.165	0.165	0.166	0.166	0.165	0.165
3rd Quintile	0.178	0.178	0.178	0.177	0.177	0.178	0.177	0.178	0.178
4th Quintile	0.188	0.188	0.189	0.187	0.187	0.188	0.187	0.188	0.188
5th Quintile	0.212	0.214	0.214	0.212	0.210	0.211	0.210	0.212	0.212

Notes:

Table A12: Share of Additional Employment Produced by Stimulus Among Workers Initially Employed (or Nonemployed) at Different Initial Earnings Quintiles (or Nonemployment):

Stimuli Consist of 500 New Positions in Alternative Combinations of Firm Size Quartile/Firm Average Pay Quartile (Averaged Across Industry Supersectors)

Earnings Quintile	Firm Size/Pay Level Combination			
	Sm./Low	Lg./Low	Sm./Hi	Lg./Hi
Nonemployed	0.439	0.453	0.412	0.428
1st Quintile	0.159	0.157	0.150	0.148
2nd Quintile	0.152	0.149	0.143	0.142
3rd Quintile	0.114	0.110	0.114	0.113
4th Quintile	0.082	0.079	0.093	0.091
5th Quintile	0.062	0.060	0.085	0.080

Notes:

Table A13: Expected Share of Utility Increases from New Stimulus Positions among Workers Initially Employed (or Nonemployed) at Different Initial Earnings Quintiles (or Nonemployment):

Stimuli Consist of 500 New Positions in Alternative Combinations of Firm Size Quartile/Firm Average Pay Quartile (Averaged Across Industry Supersectors)

Earnings Quintile	Firm Size/Pay Level Combination			
	Sm./Low	Lg./Low	Sm./Hi	Lg./Hi
Nonemployed	0.139	0.141	0.135	0.138
1st Quintile	0.120	0.120	0.117	0.118
2nd Quintile	0.167	0.167	0.164	0.164
3rd Quintile	0.179	0.178	0.177	0.177
4th Quintile	0.187	0.187	0.189	0.188
5th Quintile	0.208	0.208	0.218	0.214

Notes:

Distance from Focal Tract	Earnings Quintile					
	Nonemp.	1st Q.	2nd Q.	3rd Q.	4th Q.	5th Q.
Target Tract	6.7E-02	5.8E-03	3.7E-03	2.9E-03	2.2E-03	1.7E-03
1 Tct Away	5.4E-03	2.0E-03	1.3E-03	9.2E-04	6.5E-04	4.7E-04
2 Tcts Away	3.1E-03	1.1E-03	7.8E-04	2.3E-04	4.0E-04	3.0E-04
3+ Tcts w/in PUMA	1.6E-03	6.5E-04	4.4E-04	3.3E-04	2.4E-04	1.9E-04
1 PUMA Away	1.1E-03	4.5E-04	3.2E-04	2.3E-04	1.6E-04	1.2E-04
2 PUMAs Away	6.2E-04	2.7E-04	2.0E-04	1.4E-04	1.0E-04	7.7E-05
3+ PUMAs w/in State	4.0E-04	1.9E-04	1.3E-04	1.0E-04	7.2E-05	5.4E-05
1-2 States Away	1.2E-04	5.6E-05	3.9E-05	2.8E-05	2.0E-05	1.5E-05
3+ States Away	1.0E-05	4.8E-06	3.3E-06	2.4E-06	1.8E-06	1.4E-06

Table A14: Expected Change in Utility From New Stimulus Positions Among Workers Initially Employed at Different Combinations of Initial Earnings Quintile (or Nonemployed) and Distance from Focal Tract:
Averaged Across All Stimulus Specifications Featuring 500 New Jobs)

Distance from Focal Tract	Earnings Quintile					
	Nonemp.	1st Q.	2nd Q.	3rd Q.	4th Q.	5th Q.
Target Tract	896	363	387	440	485	545
1 Tct Away	174	208	203	214	218	234
2 Tcts Away	143	160	169	160	169	176
3+ Tcts w/in PUMA	126	152	149	157	157	165
1 PUMA Away	117	138	142	142	149	148
2 PUMAs Away	108	114	115	117	118	118
3+ PUMAs w/in State	104	109	110	112	113	113
1-2 States Away	89	92	93	93	93	94
3+ States Away	0	1	1	1	1	1

Table A15: Expected Change in Utility From New Stimulus Positions Among Workers Initially Employed in the Focal Tract at Different Earnings Quintiles (or Nonemployed) by Industry Supersector (Averaged Across Firm Size/Firm Average Earnings Combinations)

Earnings Quintile	Industry								
	Avg.	Info.	Manu.	R/W Trd.	Oth. Serv.	Ed./Hlth	Lei/Hosp.	Gov.	Const.
Nonemployed	896	702	794	860	1196	830	864	777	1142
1st Quintile	363	333	362	358	438	301	475	329	307
2nd Quintile	387	346	428	369	467	324	463	361	333
3rd Quintile	440	402	505	404	537	354	473	436	408
4th Quintile	485	453	576	451	583	347	527	477	469
5th Quintile	545	541	694	554	649	388	565	479	491

Notes:

Table A16: Expected Change in Utility From New Stimulus Positions Among Workers Initially Employed in the Focal Tract at Different Earnings Quintiles (or Nonemployed) by Firm Size Quartile/Firm Average Pay Quartile Combination (Averaged Across Industry Supersectors)

Earnings Quintile	Firm Size/Pay Level Combination			
	Sm./Low	Lg./Low	Sm./Hi	Lg./Hi
Nonemployed	1020	1051	671	842
1st Quintile	433	410	282	326
2nd Quintile	467	443	291	346
3rd Quintile	512	467	355	426
4th Quintile	507	459	449	527
5th Quintile	442	402	664	673

Notes:

Table A17: Change in Probability of Nonemployment From a Natural Disaster Removing 25, 50 or 100% of Positions in the Focal Tract for a Randomly Chosen Individual at Different Distances from Focal Tract (Averaging Across the Initial Earnings Distribution)

Distance from Focal Tract	% of Jobs Removed		
	25%	50%	100%
Target Tract	0.039	0.091	0.245
1 Tct Away	4.9E-04	8.6E-04	1.7E-03
2 Tcts Away	3.2E-04	5.5E-04	9.9E-04
3+ Tcts w/in PUMA	1.5E-04	3.0E-04	6.1E-04
1 PUMA Away	1.2E-04	2.2E-04	4.1E-04
2 PUMAs Away	7.5E-05	1.4E-04	2.8E-04
3+ PUMAs w/in State	5.3E-05	1.0E-04	2.0E-04
1-2 States Away	1.7E-05	3.4E-05	6.5E-05
3+ States Away	1.5E-06	3.1E-06	5.9E-06

Notes:

Table A18: Expected Change in Utility Produced by a Natural Disaster Removing 25, 50 or 100% of Positions in the Focal Tract Among Geographic Areas Defined by Distances from the Focal Tract: (Averaging Across the Initial Earnings Distribution)

Distance from Focal Tract	% of Jobs Removed		
	25%	50%	100%
Target Tract	-1536	- 2802	- 4653
1 Tct Away	-77	- 147	- 278
2 Tcts Away	-69	- 130	- 242
3+ Tcts w/in PUMA	-64	- 125	- 235
1 PUMA Away	-61	- 122	- 212
2 PUMAs Away	-55	- 106	- 193
3+ PUMAs w/in State	-53	- 102	- 186
1-2 States Away	-43	- 83	- 151
3+ States Away	0	- 1	- 1

Notes:

Table A19: Expected Share of Utility Decreases Produced by a Natural Disaster Among Geographic Areas Defined by Distances from the Focal Tract, by Disaster Severity (25%/50%/100% of Jobs Lost)

Distance from Focal Tract	% of Jobs Removed		
	25%	50%	100%
Target Tract	0.007	0.006	0.006
1 Tct Away	0.002	0.002	0.002
2 Tcts Away	0.004	0.004	0.004
3+ Tcts w/in PUMA	0.016	0.017	0.017
1 PUMA Away	0.023	0.023	0.022
2 PUMAs Away	0.056	0.055	0.055
3+ PUMAs w/in State	0.082	0.081	0.081
1-2 States Away	0.785	0.785	0.782
3+ States Away	0.025	0.027	0.031

Notes:

Table A20: Share of Additional Nonemployment Produced by a Natural Disaster Removing 25, 50 or 100% of Positions in the Focal Tract Among Workers at Different Initial Earnings Quintiles (or Nonemployed)

Earnings Quintile	% of Jobs Removed		
	25%	50%	100%
Nonemployed	0.330	0.315	0.284
1st Quintile	0.173	0.166	0.162
2nd Quintile	0.169	0.169	0.170
3rd Quintile	0.136	0.141	0.146
4th Quintile	0.106	0.115	0.128
5th Quintile	0.085	0.094	0.110

Notes:

Table A21: Change in Probability of Nonemployment From a Natural Disaster Removing 25% of Positions in the Focal Tract Among Workers Initially Employed at Different Combinations of Initial Earnings Quintile (or Nonemployed) and Distance from Focal Tract

Distance from Focal Tract	Earnings Quintile					
	Nonemp.	1st Q.	2nd Q.	3rd Q.	4th Q.	5th Q.
Target Tract	0.018	0.059	0.048	0.042	0.038	0.029
1 Tct Away	1.4E-03	7.5E-04	6.0E-04	3.2E-04	1.7E-04	1.5E-04
2 Tcts Away	8.0E-04	4.2E-04	2.2E-04	4.6E-04	1.0E-04	9.1E-05
3+ Tcts w/in PUMA	4.5E-04	2.8E-04	1.8E-04	9.3E-05	5.2E-05	5.4E-06
1 PUMA Away	3.0E-04	1.6E-04	1.1E-04	7.3E-05	5.4E-05	4.0E-05
2 PUMAs Away	1.8E-04	1.1E-04	7.7E-05	5.3E-05	3.8E-05	2.9E-05
3+ PUMAs w/in State	1.3E-04	7.9E-05	5.5E-05	3.7E-05	2.7E-05	2.0E-05
1-2 States Away	4.5E-05	2.2E-05	1.6E-05	1.1E-05	8.2E-06	6.2E-06
3+ States Away	4.0E-06	1.9E-06	1.3E-06	9.4E-07	6.9E-07	5.3E-07

Notes:

Table A22: Change in Probability of Nonemployment From a Natural Disaster Removing 100% of Positions in the Focal Tract Among Workers Initially Employed at Different Combinations of Initial Earnings Quintile (or Nonemployed) and Distance from Focal Tract

Distance from Focal Tract	Earnings Quintile					
	Nonemp.	1st Q.	2nd Q.	3rd Q.	4th Q.	5th Q.
Target Tract	0.052	0.304	0.281	0.280	0.279	0.239
1 Tct Away	4.2E-03	2.3E-03	1.7E-03	1.3E-03	9.7E-04	7.5E-04
2 Tcts Away	2.4E-03	1.4E-03	1.0E-03	6.1E-04	5.8E-04	4.5E-04
3+ Tcts w/in PUMA	1.5E-03	9.3E-04	6.6E-04	4.5E-04	3.1E-04	2.0E-04
1 PUMA Away	1.0E-03	5.5E-04	4.1E-04	2.8E-04	2.1E-04	1.6E-04
2 PUMAs Away	6.7E-04	3.9E-04	2.8E-04	2.0E-04	1.5E-04	1.1E-04
3+ PUMAs w/in State	4.7E-04	2.8E-04	2.0E-04	1.4E-04	1.0E-04	7.7E-05
1-2 States Away	1.7E-04	8.8E-05	6.2E-05	4.4E-05	3.3E-05	2.5E-05
3+ States Away	1.6E-05	7.8E-06	5.3E-06	3.8E-06	2.9E-06	2.2E-06

Notes:

Table A23: Expected Share of Utility Decreases Produced by a Natural Disaster among Workers Initially Employed (or Nonemployed) at Different Initial Earnings Quintiles (or Nonemployment), by Disaster Severity (25%/50%/100% of Jobs Lost)

Earnings Quintile	% of Jobs Removed		
	25%	50%	100%
Nonemployed	0.140	0.139	0.138
1st Quintile	0.120	0.120	0.120
2nd Quintile	0.166	0.166	0.166
3rd Quintile	0.178	0.178	0.178
4th Quintile	0.186	0.187	0.187
5th Quintile	0.210	0.210	0.211

Notes:

Table A24: Expected Change in Utility From a Natural Disaster Removing 25% of Positions in the Focal Tract Among Workers Initially Employed at Different Combinations of Initial Earnings Quintile (or Nonemployed) and Distance from Focal Tract

Distance from Focal Tract	Earnings Quintile					
	Nonemp.	1st Q.	2nd Q.	3rd Q.	4th Q.	5th Q.
Target Tract	-146	-1249	-1552	-1746	-1961	-2015
1 Tct Away	-67	-72	-88	-83	-73	-74
2 Tcts Away	-64	-68	-64	-76	-71	-70
3+ Tcts w/in PUMA	-62	-64	-68	-63	-63	-65
1 PUMA Away	-54	-54	-63	-61	-60	-71
2 PUMAs Away	-51	-56	-56	-54	-55	-57
3+ PUMAs w/in State	-51	-54	-54	-54	-53	-54
1-2 States Away	-41	-43	-43	-43	-43	-43
3+ States Away	0	0	0	0	0	0

Notes:

Table A25: Expected Change in Utility From a Natural Disaster Removing 100% of Positions in the Focal Tract Among Workers Initially Employed at Different Combinations of Initial Earnings Quintile (or Nonemployed) and Distance from Focal Tract

Distance from Focal Tract	Earnings Quintile					
	Nonemp.	1st Q.	2nd Q.	3rd Q.	4th Q.	5th Q.
Target Tract	-431	-3618	-4526	-5148	-6150	-6291
1 Tct Away	-227	-273	-285	-290	-287	-290
2 Tcts Away	-219	-248	-238	-252	-247	-246
3+ Tcts w/in PUMA	-211	-237	-239	-244	-243	-230
1 PUMA Away	-186	-208	-213	-225	-218	-215
2 PUMAs Away	-179	-192	-194	-195	-196	-201
3+ PUMAs w/in State	-176	-188	-187	-188	-188	-190
1-2 States Away	-145	-150	-151	-152	-152	-152
3+ States Away	0 -1	-1	-1	-1	-2	

Notes:

Table A26: Change in Probability of Destination Employment (or Nonemployment) at Different Distances from Focal Tract after a Natural Disaster Removing 25, 50 or 100% of Positions for Workers Initially Employed in the Focal Tract (Averaging Across the Initial Earnings Distribution)

Distance from Focal Tract	% of Jobs Removed		
	25%	50%	100%
Nonemployment	0.039	0.091	0.245
Target Tract	-0.102	-0.247	-0.666
1 Tct Away	0.004	0.010	0.025
2 Tcts Away	0.005	0.012	0.030
3+ Tcts w/in PUMA	0.009	0.021	0.055
1 PUMA Away	0.007	0.016	0.045
2 PUMAs Away	0.009	0.023	0.063
3+ PUMAs w/in State	0.007	0.017	0.046
1-2 States Away	0.020	0.050	0.139
3+ States Away	0.003	0.006	0.018

Notes:

Table A27: Change in Probability of Destination Employment (or Nonemployment) at Different Distances from Focal Tract after a Natural Disaster Removing 100% of Positions for Workers Initially Employed in the Focal Tract by Initial Earnings Quintile (or Nonemployment)

Distance from Focal Tract	Earnings Quintile					
	Nonemp.	1st Q.	2nd Q.	3rd Q.	4th Q.	5th Q.
Nonemployment	0.052	0.304	0.281	0.280	0.279	0.239
Target Tract	-0.068	-0.627	-0.707	-0.765	-0.804	-0.826
1 Tct Away	0.000	0.020	0.029	0.031	0.029	0.031
2 Tcts Away	0.001	0.023	0.033	0.037	0.037	0.040
3+ Tcts w/in PUMA	0.001	0.049	0.058	0.066	0.066	0.071
1 PUMA Away	0.002	0.039	0.046	0.053	0.056	0.058
2 PUMAs Away	0.003	0.055	0.062	0.072	0.077	0.087
3+ PUMAs w/in State	0.002	0.037	0.047	0.053	0.056	0.061
1-2 States Away	0.006	0.088	0.134	0.153	0.182	0.210
3+ States Away	0.000	0.012	0.016	0.019	0.023	0.029

Notes:

Table A28: Change in Probability of Destination Employment (or Nonemployment) at Different Distances from Focal Tract after a Natural Disaster Removing 25% of Positions for Workers Initially Employed in the Focal Tract by Initial Earnings Quintile (or Nonemployment)

Distance from Focal Tract	Earnings Quintile					
	Nonemp.	1st Q.	2nd Q.	3rd Q.	4th Q.	5th Q.
Nonemployment	0.018	0.059	0.048	0.042	0.038	0.029
Target Tract	-0.025	-0.123	-0.122	-0.117	-0.109	-0.104
1 Tct Away	0.000	0.004	0.005	0.005	0.004	0.004
2 Tcts Away	0.000	0.005	0.006	0.006	0.005	0.006
3+ Tcts w/in PUMA	0.001	0.010	0.011	0.011	0.010	0.009
1 PUMA Away	0.001	0.008	0.008	0.008	0.007	0.007
2 PUMAs Away	0.001	0.011	0.011	0.011	0.010	0.011
3+ PUMAs w/in State	0.001	0.007	0.008	0.008	0.008	0.008
1-2 States Away	0.002	0.017	0.023	0.023	0.024	0.026
3+ States Away	0.000	0.002	0.003	0.003	0.003	0.003

Notes: