

WORKING PAPER · NO. 2021-46

Productivity Growth and Workers' Job Transitions: Evidence from Censal Microdata

Elias Albagli, Mario Canales, Chad Syverson, Matias Tapia, and Juan Wlasiuk

APRIL 2021

Productivity Growth and Workers' Job Transitions: Evidence from Censal Microdata*

Elías Albagli Mario Canales Chad Syverson
Central Bank of Chile Central Bank of Chile Chicago Booth and NBER

Matías Tapia Juan Wlasiuk
Central Bank of Chile Central Bank of Chile

March 30, 2021

Abstract

A large body of work has highlighted the importance of employment reallocation as a driver of aggregate productivity growth, but there is little direct evidence on the extent and nature of this process at the worker-firm level. We use an administrative matched employer-employee census for Chile to provide novel insights into the relationship between job transitions and productivity variation across firms. As many theories would predict, worker flows from lower- to higher-productivity firms are larger than those of the opposite sign. Empirically, however, this is only marginally so. Almost half of all transitions occur “down the firm productivity ladder.” This process is also highly heterogeneous along several dimensions. Up-the-ladder flows are more likely for direct job-to-job transitions than those that pass through non-employment. They are also much more likely for young, high-skilled workers, whose job transitions comprise in an accounting sense the lion’s share of aggregate productivity change. Interestingly, workers with the highest job turnover rates contribute proportionally the least to aggregate productivity changes. Put together, this evidence implies that the productivity mechanics of job reallocation yields a net benefit, but this hides massive and heterogeneous gross flows underneath.

JEL Classification: D2, E2, J6.

Keywords: Productivity, reallocation, job transitions, heterogeneity .

*We thank seminar participants at the Central Bank of Chile Workshop on Labor Markets and Productivity. The views expressed in this paper are exclusively those of the authors and do not necessarily reflect the position of the Central Bank of Chile or its Board members. Any errors or omissions are responsibility of the authors. Emails: ealbagli@bcentral.cl, mcanales@bcentral.cl, mtapia@bcentral.cl, chad.syverson@chicagobooth.edu, jwlasuik@bcentral.cl.

1 Motivation

Labor markets reallocate workers across firms and sectors over time. The ease with which such reallocation responds to shifts in supply and demand primitives at the micro level has considerable influence on businesses' and industries' trajectories, frictional unemployment, and over longer horizons, aggregate productivity growth. On this last point, a large body of work suggests that reallocation of labor and other factors works in the “right direction”—that is, away from firms with lower marginal products and toward those with higher marginal products. In aggregate, this process explains a relevant share of an economy's productivity gains.

At the same time, however, substantial differences in measured productivity across firms persist. Frictions limit the speed and completeness of productivity-enhancing reallocations.

In this paper we seek to make some headway in understanding how labor market reallocation interacts with the presence of heterogeneous-productivity firms. We both quantify the influence of labor reallocation on aggregate productivity growth and learn more about the frictions that limit its influence.

We conduct our investigation by leveraging a matched employer-employee census for Chile between 2005 and 2016. Chile is an interesting case study in labor reallocation dynamics for several reasons. It is a small open economy with substantial exposure to relative price shocks, particularly in commodities. This, alongside the influence of technological change, leads to a significant amount of labor reallocation across firms and sectors. In fact, Chile ranks highest among OECD countries in terms of labor turnover, with large job creation and destruction rates at the firm level and short employment spells for individual workers (Albagli et al., 2017). At the same time, as in many other countries, there is a large degree of persistent measured misallocation across Chilean firms, pointing to sizeable potential productivity gains through employment reallocation (Busso et al., 2013; Bergoing et al., 2010). This poses a puzzle directly tied to our core question: if the labor market is so fluid, why are productivity differentials across firms so persistent? Do individual job transitions really correlate with productivity differentials?

Our data allow us to track individual workers' job histories at monthly frequency and to compute annual firm-level productivity measures across all economic sectors. This enables a detailed look at how productivity gaps across firms relate to worker flows among them, as well as a direct quantification of the relative contributions of different worker groups to aggregate productivity gains from reallocation.

We highlight three main results. First, while on average individual job transitions move workers from lower-productivity to higher-productivity firms, the fraction of all job changes that flow up the productivity ladder is only marginally higher than those in the opposite direction. Almost half (49% in our preferred specification) of all job changes move workers towards firms that have lower productivity levels than their prior employer. Thus net labor reallocation towards higher-productivity firms hides a very large degree of job churning whose productivity implications mostly cancel.

Second, we show that the fraction of up-the-productivity-ladder transitions varies significantly across job flow types and across firm productivity levels. They are more likely for direct job-to-job transitions than for those passing thorough non-employment spells. This is consistent with standard job-search theory, in the sense that movements resulting from on-the-job search should lead to better quality jobs more likely to be available at higher productivity firms. They are also more likely to occur between firms at the high end of the productivity distribution, with up-the-ladder transitions originating from low productivity firms happening less frequently than implied by a model of completely random movements.

Third, the job transitions of younger workers, workers with high skills, and female workers are more likely to be to higher-productivity firms and correspond to larger productivity gains. In fact, young skilled workers provide the plurality of net labor reallocation to higher productivity firms. The net contributions of other large groupings are modest or even negative. Workers with the highest turnover rates contribute proportionally little to aggregate productivity growth.

Overall, our findings sound a note of caution, or at least modesty in interpretation, about the connections between labor market fluidity and aggregate productivity growth. While fluidity enables reallocation of resources to more productive firms, many of the observed job transitions appear to happen for reasons that are orthogonal to industry or aggregate productivity growth. The flip side of this coin is that there may be a substantial amount of untapped potential for labor reallocations to enhance productivity growth.

Our paper relates to several strands of literature. The first are studies documenting and analyzing the existence of persistent productivity differences across firms, and the potential efficiency gains of unrealized labor reallocation (Hsieh and Klenow, 2009; Restuccia and Rogerson, 2013; Dias et al., 2016). A second literature looks more directly at the role of actual employment reallocation across firms as a relevant driver of aggregate productivity growth (De Loecker and Konings, 2006; Foster et al., 2008; Syverson, 2011; Melitz and Polanec, 2015). These papers are closest to the spirit of our analysis. We add to their insights by directly quantifying flows up versus down the productivity ladder, as well as the actual contributions to aggregate efficiency gains of different types of job-flows and worker types. A third body of evidence studies the determinants and implications of individual job transitions (Topel and Ward, 1992; Fallick et al., 2011; Davis et al., 2012; Foster et al., 2016; Sorkin, 2018) and, more recently, links such transitions to movements along job ladders from the perspective of workers (Menzio et al., 2016; Lise and Robin, 2017; Moscarini and Postel-Vinay, 2018; Bagger and Lentz, 2019). Our results are generally consistent with their findings, but complement them by linking evidence from the point of view of workers to that of productivity gaps between firms.

We describe our data and methodology in the next section. Section 3 presents evidence of overall efficiency gains from job flows, differentiating by job transition type and by firm of origin productivity levels. Section 4 further studies the role of heterogeneity at the worker level, and provides a direct quantification of the relative contribution of different groups to aggregate productivity-gains. Section 5 concludes.

2 Data and Methodology

Our data comes from two main sources. First, the Chilean tax authority, SII, provides tax records for the universe of Chilean firms from 2005 through 2016. The annual tax statement (F22) provides information on each firm’s balance sheet, including data on sales, expenditures in intermediate inputs, the wage bill, and capital stocks. Additionally, firms file a labor compensation report (DJ1887) that records the annual compensation paid to each employee and the specific months in which she was employed at the firm. We use this information to build a monthly matched employer-employee dataset in which we observe the employment status of an individual worker, earnings, and any job transitions. Because all formal firms must report to the SII, the data covers any worker with a formal wage contract. This represents roughly 60% of all employment in Chile. All data are anonymized to ensure confidentiality of firms’ and workers’ identities.

Second, we obtain workers’ genders and dates of birth from the Chilean Civil Registry and Identification Service (Servicio de Registro Civil e Identificación) and match it to the SII data.

We apply a set of filters to obtain the final dataset in our empirical analysis. First, we drop all firms that have only one employee for an entire year, as it seems likely that they represent some form of self-employment rather than an actual firm with workers in the spirit of our exercise. We also drop all workers employed in those firms. This is about 1.7% of the initial set of 48 million monthly worker-job matches. Second, we want to avoid spurious job transitions in which workers move to a “new” firm with a different tax ID that is actually directly linked to the previous firm. This includes cases in which the firm tax ID changes, M&As, or separations of a single firm into several business entities for tax reasons. We address this by excluding from the set of transitions all cases in which a significant share of a firm’s workforce jointly reallocates to another firm, whether existing or new. This removes just under 5000 firms.¹ Third, we must assign any workers employed by more than one firm in a given month to a single job and employer. We classify the worker’s main job as the one with the longest current tenure and, if there is a tie, the one with the highest average monthly earnings. About 8% of our worker-job observations are dropped as secondary jobs. Finally, as we want to focus on full-time jobs but have no information on hours, we drop all employment relations in which the worker earns an implied full-time wage that is less than the minimum wage for 80% of her tenure. These jobs, which account for 20% of the initial set of observations, are probably part-time.

We calculate average labor productivity at the firm level as a measure of annual sales (from form F22) per annual-equivalent worker. Firms’ annual-equivalent workers are simply the sum of workers in firms’ monthly DJ1887 reports, divided by 12. For example, a firm with one worker who is employed the entire year and two employees who each work in six months of the year (whether overlapping months or not) has two annual-equivalent

¹We use a sliding scale to define *en masse* transitions likely to reflect spurious firm changes. Namely, we do not consider job transitions where the following fractions of workers in a firm of a given size category jointly move to the same firm: 100% of workers in firms with 4-5 employees, 90% of workers in firms with 6-10 employees, 80% of workers in firms with 11-20 employees, 70% in firms with 21-50 employees, 60% in firms with 51-100 employees, 50% in firms with 101-200 employees, or 20% in firms with over 200 employees. These thresholds are of course rather arbitrary, but results were not affected when choosing different reasonable thresholds.

employees.² Thus, average labor productivity is defined at an annual level.³

We define job transitions as changes in the firm identity tied to the worker’s main job. If this occurs over two consecutive months, we consider it a job-to-job transition. If there is a gap of a month or more between jobs, we label it as a transition through a non-employment spell. This could reflect unemployment, exit from the labor force, or given the limits of our data, a spell of employment in the informal sector.

We do not have worker education data, so we use workers’ earnings to classify them into skill groups. We follow the methodology proposed by [Borovičková and Shimer \(2017\)](#), who define a worker’s skill type as the expected wage she should receive in an employment relationship, conditional on accepting the job. This provides a measure of the market’s valuation of the worker’s skills, reflecting human capital acquired through both education and experience. Results throughout the paper are robust to simpler methodologies, such as a worker’s earnings quantile within gender-cohort groups.

To compute the [Borovičková and Shimer \(2017\)](#) skill group, we first obtain the residual from a regression of logged monthly real earnings of worker i in firm j and year t , w_{ijt} , on indicators for the year, d_t , as well as for the worker’s age-by-gender group, $d_{a,g}$. The regression is shown in equation 1.

$$w_{ijt} = \beta_0 + \beta \sum_{t=2005}^{2016} d_t + \delta \sum_{a,g} d_{a,g} + \varepsilon_{ijt} \quad (1)$$

We average this residual through each worker’s observed employment history to obtain an estimate of the worker’s type. Finally, we rank workers by type into quintiles, weighting by the worker’s observed length of employment. Because the methodology removes the time component of wages, this worker ranking is time-invariant and comparable across the sample.

We also classify workers based on their job turnover propensity. Using observed job tenure histories, we run Cox survival regressions on the duration of employment spells, including those that remain open at the end of the sample. These hazard models control for gender and age. The residuals are ranked and workers sorted into separation hazard quintiles. These residuals reflect an estimate of an individual worker’s idiosyncratic propensity toward high or low job turnover.

²For consistency with the concept of labor productivity, we include all workers employed in the firm at any given month, regardless of whether they had multiple jobs.

³For robustness, we repeated all our analyses below using alternative measures of productivity, such as TFPR (actually, here, the residual of the revenue function) and the marginal revenue product of labor constructed using the methodology in [Akerberg et al. \(2015\)](#). We do not present these results here for the sake of brevity, but they are qualitatively similar and available upon request.

3 Job Transitions and Productivity Sorting Patterns

We begin our analysis by relating worker transitions between firms to differences in those firms' labor productivity levels.

For each transition, we calculate the difference between the logged labor productivity levels of the destination and origin firms. This measures both the direction of any given transition—whether the worker went up or down the productivity job ladder—as well as the magnitude of the change. We present statistics of this distribution in Table 1.

Table 1: Job Transitions and Productivity Gaps

	Average Productivity Gap (1)	p25 (2)	p50 (3)	p75 (4)	Share Upward Trans. (5)	Transi- tions (6)
<i>Panel A: Labor Productivity</i>						
All	7.6	-62.5	5.2	76.1	52.4	11,017,590
Job-to-Job	15.9	-51.4	8.3	78.2	54.2	4,160,740
Job-N-Job	2.6	-69.7	2.9	74.9	51.2	6,856,850
<i>Panel B: Adjusted Labor Productivity</i>						
All	4.6	-63.6	2.8	71.6	51.3	11,017,590
Job-to-Job	14.0	-51.6	7.5	75.7	53.8	4,160,740
Job-N-Job	-1.0	-71.1	-0.5	69.2	49.8	6,856,850

Note: Productivity measured as log average labor productivity. Productivity gaps are defined as the difference between the firm of destination and the firm of origin for any given job transition. Values multiplied by 100 for readability. “Upward transitions” are defined as job transitions with a positive productivity gap.

Source: Authors' calculations based on Chilean SII data.

Panel A presents the results across all 11 million transitions in our sample. Consistent with reallocation being a factor behind productivity growth, the average labor productivity gap between the destination and origin firm in a worker transition is about 8%. However, this aggregate result averages over enormous heterogeneity. This is clear in the large dispersion of the productivity differentials. The 25th percentile of the distribution involves a worker moving to a firm with a labor productivity level about half the size ($e^{-0.625} = 0.535$) of the firm she left. The 75th percentile transition is to a firm with a productivity level more than twice as high ($e^{0.761} = 2.14$). The rough symmetry of this dispersion is reflected in the fact that the share of job transitions that move up the productivity ladder is only a shade above half, with 47.6% of worker transitions being to firms less productive than the origin firm.

This is a striking result. While net reallocation flows are directed towards more productive firms, the process involves an enormous degree of churning in both directions. Theoretically, even in a frictionless world in which all job transitions were efficient, some

transitions down the firm-productivity ladder might be motivated by better firm-worker matches or from other non-pecuniary motives that might drive worker reallocation across different firms (Sorkin, 2018). Nonetheless, the fact that the frequency of upward flows is only marginally higher than that of a completely random reallocation benchmark, with a 50% unconditional change of moving up or down the firm-productivity ladder, is notable.

To our knowledge, there is little comparable evidence available in the prior literature. A notable exception is the work by Stoyanov and Zubanov (2012), who study labor flows and productivity in Denmark’s manufacturing sector. Their research focus is different from ours, centering on the potential externalities created when workers from more to less productive firms, rather than a comprehensive characterization of reallocation and productivity. However, they do mention the share of upward transitions in their data is 55%. This figure is somewhat higher but still close to our findings.

The second and third rows of Panel A show the separate distributions for direct job-to-job transitions and those broken by a spell outside formal employment (job-N-job). Job-to-job transitions account for 37% of all transitions in our sample. They are associated with larger average productivity gaps between destination and origin firms, and indeed the whole distribution of productivity differentials shifts to the right. Workers that change jobs directly are therefore more likely to climb productivity ladders and their movements correspond to larger productivity differentials. In fact, if we sum all productivity changes tied to every reallocated worker, job-to-job transitions account for almost 80% of this implied aggregate productivity gain.

This pattern is consistent with the presence of search frictions in the reallocation process. The idea of on-the-job search introduced by Burdett and Mortensen (1998) implies that direct job-to-job transitions are likely to represent voluntary movements to better matches. To the extent that match qualities correlate positively with measures of firm-level productivity, there should be a link between direct job-to-job transitions and worker moving upwards through a productivity ladder. On the other hand, the productivity-enhancing effects of indirect job transitions, where the worker goes through an intervening spell without formal employment, are plausibly weaker. For example, a worker who experienced an involuntary separation from her prior employer and had to find work might end up accepting a worse match, as long as it satisfies her reservation wage.

All this said, while the share of upward transitions is larger among job-to-job transitions, the fraction is still relatively close to 50%. Even direct job-to-job transitions reflect a very noisy reallocation process.

Panel B addresses a potential concern regarding the exercise shown in Panel A. As discussed earlier, our data defines firm productivity by calendar years. Transitions that occur between calendar years compare productivity measures across different points in time. Secular aggregate productivity growth implies that the average true productivity gaps for across-year transitions would be overestimated. Additionally, the raw productivity gap does not account for systematic differences in average labor productivity across sectors, which are likely related to differences in capital intensity. We therefore adjust productivity across years by subtracting from the measured destination-origin productivity difference both the average annual productivity growth over the period as well as, for

transitions across sectors, the difference in mean sector productivity levels.

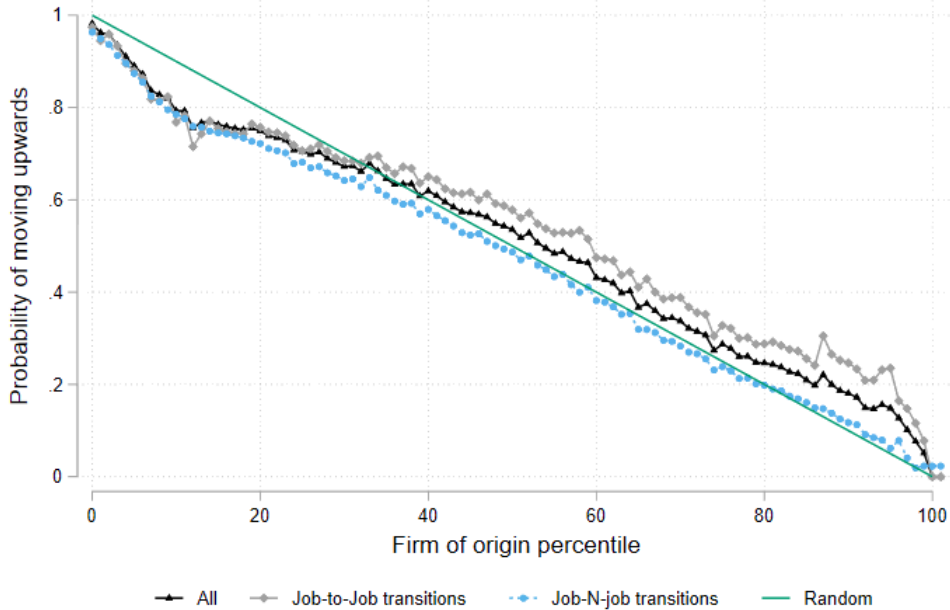
After this adjustment, the average productivity gaps do become smaller, with the average differential for all transitions falling to about 5%, and 51.3% of transitions being upwards. Most of the basic patterns in Panel A remain, however. The biggest difference is that the average and median productivity gaps of job-N-job transitions are now slightly negative. For the rest of the paper, our attention will focus on this measure of adjusted labor productivity differences across transitions, although all results are qualitatively robust to using the unadjusted gaps summarized in Panel A.

To have a sense of the macroeconomic scale of the reallocation process described by the data, we sum sales per worker gaps across all job transitions that take place during a given year and then divide that sum by firms' total sales that year. This calculation assumes each transition leads to an output change equal to what would happen if the transitioning worker experienced a labor productivity change equal to that between of the two firms she transitions between. The sum of these output changes relative to total output provides a metric of output growth coming from productivity gains through labor reallocation. Over our sample period, the implied average annual sales growth from reallocation computed in this way is 1.35%. For comparison, average annual total sales growth in our data is 5.9%. This suggests that labor reallocation from lower to higher marginal product activities is quantitatively important as a potential source of growth.

Figure 1 presents another way to summarize the connections between labor reallocation and productivity. It plots, for every percentile of the firm productivity distribution, the probability that a given worker transition from an origin firm at that percentile is to a destination firm with a higher productivity level. If worker reallocations were entirely unrelated to productivity differentials, the probability would be given by a negative 45-degree line. For example, a transitioning worker leaving an origin firm at the 25th percentile would have a 75 percent probability of moving to a higher-productivity firm. As the figure shows, actual reallocation patterns depart from the random benchmark, but modestly so, and the deviation from randomness varies systematically through the productivity distribution. Movements up the productivity ladder are disproportionately likely among workers employed by firms in the upper half of the productivity distribution. In contrast, workers leaving firms at the lower tail of the productivity distribution are more likely to move to an even lower-productivity firm than if they moved randomly.

It is also interesting that at the bottom end of the productivity distribution, job-to-job and indirect transitions look virtually identical. For workers at low-productivity firms, there seems to be no systematic difference between changing jobs directly or indirectly in terms of the likelihood of moving up the productivity ladder. In contrast, the two types of transitions look very different at the upper end of the productivity distribution. Indirect transitions lie (almost) along the random reallocation benchmark, but job-to-job transitions are much more likely to lead to movements up the job ladder.

Figure 1: Conditional probability of moving to a more productive firm



Note: Job transitions are sorted by adjusted labor productivity of the firm of origin. For each percentile of this distribution, the figure plots the share of transitions to a more productive firm.
 Source: Authors’ calculations based on Chilean SII data.

4 Job Transitions and Worker Heterogeneity

In this section, we explore systematic differences in the labor reallocation-productivity relationship across observable worker types.

4.1 Reallocation Patterns by Worker Groups: Firm-Level Evidence

For a first glance at the role of worker heterogeneity in employment transitions and productivity gaps, we replicate the exercise in Haltiwanger et al. (2018), who look at job-to-job flows between low and high productivity firms (what they call “poaching flows”) in the US across different worker categories. This provides an analysis at the firm level, rather than at the worker level as the rest of the paper, but provides some insights that can motivate our later analysis.

Haltiwanger et al. (2018) place firms into productivity quintiles and focus on net job-to-job flows at the lowest and highest quintiles. Like our findings above, their results show that net job-to-job flows are consistent with productivity-enhancing reallocation. Net flows are negative for low-productivity firms and positive for high-productivity firms. However, net inflows into high-productivity firms are not homogeneous across types of individuals. Younger and more educated workers are over-represented.

We find very similar patterns in Chile. Net job-to-job flows are negative (positive) for

Table 2: Poaching Patterns by Worker Type

	Share in the economy (1)	Low Productivity		High Productivity	
		Net Job-to-job (2)	Ratio (3)	Net Job-to-job (4)	Ratio (5)
<i>Panel A: By age</i>					
Less than 25	13.0	26.7	2.05	14.2	1.09
Age 25 - 34	31.1	46.6	1.50	48.1	1.55
Age 35 - 44	26.5	19.4	0.73	24.3	0.92
Age 45 or older	29.4	7.3	0.25	13.4	0.45
<i>Panel B: By worker skill type</i>					
Low wage	23.4	10.0	0.43	5.8	0.25
Q2	23.3	24.7	1.06	15.7	0.67
Q3	20.6	29.5	1.43	22.9	1.11
Q4	17.3	25.8	1.49	28.9	1.67
High wage	15.4	10.0	0.65	26.7	1.74

Note: This table summarizes net job-to-job transitions by age and worker skill categories. Following [Haltiwanger et al. \(2018\)](#), we define high- (low-)productivity firms as the top (bottom) quintile in the firm labor productivity distribution.

Source: Authors' calculations based on Chilean SII data.

low (high) productivity firms across all worker categories. This reinforces the idea that on net, reallocation enhances productivity. More interestingly, though, poaching is far from uniform across all workers. This can be seen in Table 2.

Panel A presents poaching patterns across four age groups: workers younger than 25 years old, 25-34 years old, 35-44 years old, and 45 and older. Column (1) presents the worker age groups' shares of all job-to-job transitions in the sample, while column (2) reports the groups' shares of job-to-job transitions out of firms in the lowest productivity quintile. As is clear in column (3), which reports the ratio of these two shares, workers poached from these low-productivity firms are considerably younger. Even though those under 35 account for only 44% of transitions, they are 73% of poached workers. Workers over 45, while almost 30% of transitions, represent a scanty 7% of workers who leave low-productivity firms. This is not completely a phenomenon of low-productivity firms. If we focus on poaching from firms in the highest productivity quintile, as shown in columns (4) and (5), the age pattern is not as strong yet still remains. More than 60% of workers poached from these firms are younger than 35, and those over 45 are underrepresented.

Panel B repeats the exercise for worker skill groups, where worker skills are proxied by the lifetime earnings quintiles defined in Section 2. Once again, poaching patterns are far from homogeneous. Here, the strongest patterns regard transitions away from the highest-productivity firms. These are disproportionately high-skill (high-wage) work-

ers, with the highest skill quintile accounting for 1.7 times as many transitions from the most productive firms as their overall share of transitions. Poaching from the lowest-productivity firms is qualitatively similar but of smaller magnitude.

Overall, the results in Table 2 indicate that worker heterogeneity is tied to systematic differences in reallocation patterns. These differences are consistent with existent theory and evidence on labor market dynamics, and are similar to the evidence found for the US in Haltiwanger et al. (2018).

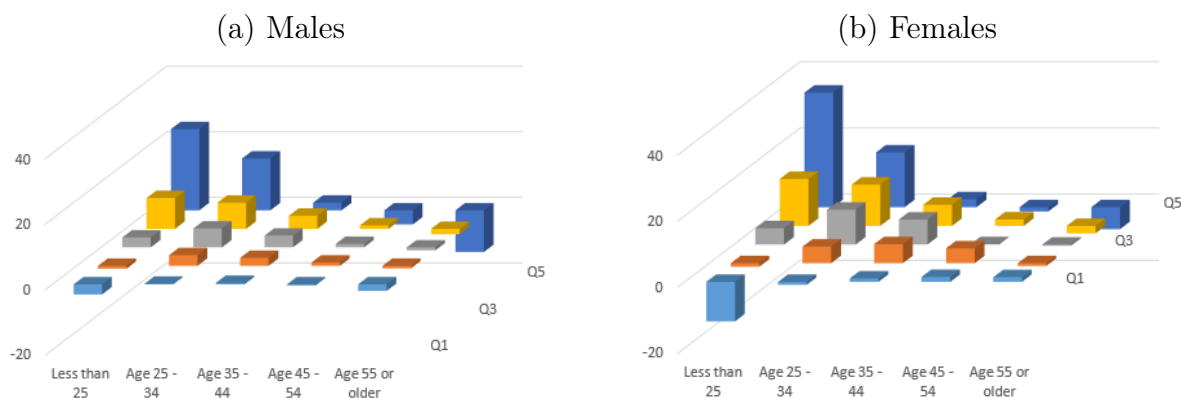
4.2 Reallocation Patterns by Worker Groups: Worker-Level Evidence

An enormous literature (Topel and Ward, 1992; Menzio et al., 2016; Lagakos et al., 2018) has documented the importance of worker life-cycles in shaping employment dynamics and the evolution of earnings. Job mobility is larger early in the employment life-cycle, consistent with both young workers starting to build their job ladders and having lower transition costs due to having built less sector- or firm-specific human capital. Earnings growth is also especially strong in the first decade of labor market participation, pointing to a rapid accumulation of skills and rapidly movement to better job matches. Workers in the later portions of their work lives, on the other hand, move significantly less, have flatter earnings profiles, and suffer significantly more from job displacements. It is therefore not surprising that younger workers play a leading role in employment reallocation towards more productive firms.

Similarly, there are reasons why movements up the productivity ladder might be more likely for workers in the upper end of the ability distribution. For example, complementarities between high-productivity firms and high-skilled workers will naturally lead to sorting as an equilibrium outcome.

To explore these elements further and complement our firm-heterogeneity results in the previous section with worker-heterogeneity results, we place every worker into one of 25 age-by-skill worker groups. These use the five skills quintiles derived as above as well as five age groups (less than 25, 25-34, 35-44, 45-54, 55+). We compute the average productivity gap associated with worker transitions within each group. We take further advantage of our data to introduce one additional dimension by dividing the sample into men and women. This is a particularly interesting dimension for an economy like Chile, in which the gender wage gap is large and female labor participation, while growing, is still low.

Figure 2: Productivity gaps by worker characteristics



Note: The figure plots adjusted average productivity gaps resulting from job transitions, averaged by worker type group.

Source: Authors' calculations based on Chilean SII data.

Figure 2 graphs the results. The differences across groups reflect some remarkable patterns.

Consistent with the discussion above, productivity differentials are significantly larger for young, skilled workers. They are the largest for workers younger than 25 and in the top skill quintile. Productivity gaps decline monotonically with age, even conditioning on skill, becoming almost negligible for workers over 45. They also tend to increase with worker skill, conditioning on age. The one group where the average productivity difference between the destination and origin firm is clearly negative is for high-skilled workers over 55. This may reflect that a large share of those transitions are non-voluntary and associated with the destruction of such workers' valuable job ladders. Comparing across genders, productivity gaps for a given age-skill group are consistently larger for female workers. The sole exception is the youngest, least-skilled group, where long time gaps between jobs are frequent. This indicates systematic gender-related difference in the process of labor reallocation.

Overall, these results are further evidence of the heterogeneity in the interaction between job transitions and productivity. Productivity-enhancing reallocation is significantly stronger among young, high-skilled, and female workers. We now use these results to decompose productivity gains from reallocation into group-level contributions.

4.3 Taking Stock: Decomposing Aggregate Gains from Job Reallocation across Different Groups

We explore the aggregate implications of the heterogeneity documented above by computing the total implied productivity change from worker transitions within each group. For each worker category, we calculate the contribution as the average productivity differential from job transitions within the group, weighted by the group's share of total transitions. As with the analogous exercise we did above for the aggregate, this assumes that each transitioning worker experiences a labor productivity change equal to the pro-

ductivity gap (positive or negative) between her destination and origin firm. To the best of our knowledge, this decomposition of aggregate reallocation across worker characteristics is novel to the literature, and provides a quantitative assessment of the microeconomic structure that underlies macroeconomic reallocation.

Table 3 presents the results. Column (1) decomposes the economy’s average productivity gain from job transitions (4.69%) into the absolute contributions by age group (Panel A), worker skill quintile (Panel B), gender (Panel C), and job tenure (turnover propensity) quintiles (Panel D). Column (2) reports the share of total reallocation accounted for by each group, while Column (3) reports the share of job transitions. Column (4) is the ratio of (2) to (3).

As suggested by the analysis in Figure 2, there is a lot of heterogeneity in implied productivity effects across different types of workers. Almost half of the economy’s worker-reallocation-based productivity gains are associated with the job transitions of workers between 25 and 34 years old, and this rises to 80% if those under 25 are included. Productivity-enhancing labor reallocation is a young worker’s game. Job transitions of workers between 45 and 54 contribute a meager 0.8%, and those over 55 years old make a negative contribution. This once again reflects the strong life-cycle effect in reallocation patterns.

Column (4) highlights that these results reflect more than the fact that workers under 35 are more mobile. Their share of the implied productivity gain considerably exceeds their share of transitions. Younger workers not only transition more, they experience a larger productivity gap when they do move.⁴ As discussed above, these age-specific reallocation patterns closely relate to the literature on life-cycle income profiles (Lagakos et al., 2018; Menzio et al., 2016), which highlights that wage gains—likely associated with successful job transitions towards better matches—are especially relevant during the first decade of work experience.

The worker skill decomposition in Panel B shows that average productivity differentials are non-negative for all quintiles, but the relative contribution toward aggregate productivity growth rises monotonically with our skills measure. Transitions from workers in the first quintile make no contribution to productivity growth through reallocation. Conversely, almost 60% of all reallocation gains come from the top two skill quintiles despite representing only 38% of transitions. In fact, although workers in the highest quintile move relatively infrequently (accounting only for 16% of all transitions), they explain more than a third of aggregate reallocation-based labor productivity gains. These results highlight that, conditional on switching jobs, more skilled individuals are more likely to move to a firm with larger labor productivity. Productivity job ladders among more skilled individuals are a relevant driver of reallocation-based productivity gains in the economy.

The third panel replicates the exercise for male and female workers. Female workers punch above their weight in terms of productivity gains through labor reallocation. While they account for only 21% of job transitions in the economy, they comprise 31% of reallocation-based aggregate productivity gains.

⁴As mentioned earlier, they are also more likely to move directly through job-to-job transitions.

Table 3: Decomposition of Average Productivity Gains from Reallocation

	Weighted average productivity gains	Share of productivity gains	Share of job transitions	Gains-to- transitions ratio
	(1)	(2)	(3)	(4)
<i>Panel A: By age</i>				
Less than 25	1.46	31.0%	18.3%	1.69
Age 25 - 34	2.66	56.7%	35.3%	1.61
Age 35 - 44	0.70	14.9%	24.8%	0.60
Age 45 - 54	0.04	0.8%	15.4%	0.05
Age 55 or older	-0.16	-3.4%	6.2%	-0.55
Total	4.69			
<i>Panel B: By worker skill quintile</i>				
Low skill	0.00	0.0%	15.4%	0.00
Q2	0.56	12.0%	24.3%	0.49
Q3	0.98	21.0%	22.8%	0.92
Q4	1.58	23.6%	21.6%	1.56
High skill	1.57	33.5%	16.0%	2.09
Total	4.69			
<i>Panel C: By gender</i>				
Male	3.25	69.1%	79.3%	0.87
Female	1.45	30.9%	20.7%	1.49
Total	4.69			
<i>Panel D: By job tenures (turnover propensity)</i>				
Short tenure	0.59	12.5%	38.1%	0.33
Q2	0.92	19.5%	25.9%	0.75
Q3	1.03	22.0%	17.8%	1.24
Q4	1.17	24.9%	11.7%	2.13
Long tenure	0.99	21.1%	6.6%	3.2
Total	4.69			

Note: Productivity measured as log average labor productivity. Weighted average productivity gains are the adjusted average productivity gap in a group's worker transitions weighted by that group's share of total transitions.

Source: Authors' calculations based on Chilean SII data.

Panel D looks at the contribution of workers with different job tenures, based on the job turnover propensity regressions described in Section 2. Workers in the two quintiles with the shortest employment spells have a high turnover rate, and account for 64% of all transitions. However, they represent only 32% of total reallocation gains. The bulk of reallocation-based productivity gains come from quintiles 3 and 4, which contribute 47% of all gains despite being only 29% of job transitions. Workers on the fifth quintile by definition seldom change jobs, and therefore represent a meager 6.6% of job transitions in the economy. Despite this, when they do move, they experience very large productivity gains, providing 21% of the overall productivity growth through worker reallocation. This general pattern is seen in the gains-to-transitions ratio in column (4), which increases strongly along the tenure groups and even more so than across skill levels.

5 Conclusions

This paper provides a novel and more complete characterization of the relationship between individual job transitions and the aggregate process of employment reallocation that increases productivity by moving workers towards more productive firms.

Our results suggest that this process has a complex structure. The labor market's ability to reallocate workers away from less productive and into more productive firms involves an enormous amount of labor turnover, with a very large share of job transitions not leading to net productivity gains. Productivity-enhancing job transitions are not uniformly distributed, but instead exhibit significant and systematic differences across the distribution of firms and workers. Such transitions draw especially heavily from young, high-skill workers. Workers who change jobs the most frequently contribute proportionally the least to reallocation-based productivity growth. These patterns of heterogeneity, besides being informative about some theories of the labor market and serving as useful quantitative benchmarks, also demonstrate that a highly fluid labor market need not be an unequivocal sign of economy-wide productivity gains.

References

- Akerberg, Daniel A, Kevin Caves, and Garth Frazer, “Identification properties of recent production function estimators,” *Econometrica*, 2015, 83 (6), 2411–2451.
- Albagli, Elías, Alejandra Chovar, Emiliano Luttini, Carlos Madeira, Alberto Naudon, and Matias Tapia, “Labor Market Flows: Evidence from Chile Using Micro Data from Administrative Tax Records,” *Working Papers Central Bank of Chile*, 2017, 812.
- Bagger, Jesper and Rasmus Lentz, “An empirical model of wage dispersion with sorting,” *The Review of Economic Studies*, 2019, 86 (1), 153–190.
- Bergoing, Raphael, Facundo Piguillem, and Norman V Loayza, *Why are developing countries so slow in adopting new technologies? The aggregate and complementary impact of micro distortions*, The World Bank, 2010.
- Borovičková, Katarína and Robert Shimer, “High wage workers work for high wage firms,” Technical Report, National Bureau of Economic Research 2017.
- Burdett, Kenneth and Dale T Mortensen, “Wage differentials, employer size, and unemployment,” *International Economic Review*, 1998, pp. 257–273.
- Busso, Matías, Lucia Madrigal, and Carmen Pagés, “Productivity and resource misallocation in Latin America,” *The BE Journal of Macroeconomics*, 2013, 13 (1), 903–932.
- Davis, Steven J, R Jason Faberman, and John Haltiwanger, “Labor market flows in the cross section and over time,” *Journal of Monetary Economics*, 2012, 59 (1), 1–18.
- Dias, Daniel A, Carlos Robalo Marques, and Christine Richmond, “Misallocation and productivity in the lead up to the Eurozone crisis,” *Journal of Macroeconomics*, 2016, 49, 46–70.
- Fallick, Bruce, John Haltiwanger, and Erika McEntarfer, “Nonemployment Duration and the Consequences of Job Separations,” *Unpublished*. Accessed November, 2011, 15, 2011.
- Foster, Lucia, Cheryl Grim, and John Haltiwanger, “Reallocation in the great recession: cleansing or not?,” *Journal of Labor Economics*, 2016, 34 (S1), S293–S331.
- , John Haltiwanger, and Chad Syverson, “Reallocation, firm turnover, and efficiency: Selection on productivity or profitability?,” *American Economic Review*, 2008, 98 (1), 394–425.
- Haltiwanger, John, Henry Hyatt, and Erika McEntarfer, “Who moves up the job ladder?,” *Journal of Labor Economics*, 2018, 36 (S1), S301–S336.
- Hsieh, Chang-Tai and Peter J Klenow, “Misallocation and manufacturing TFP in China and India,” *The Quarterly Journal of Economics*, 2009, 124 (4), 1403–1448.
- Lagakos, David, Benjamin Moll, Tommaso Porzio, Nancy Qian, and Todd Schoellman, “Life cycle wage growth across countries,” *Journal of Political Economy*, 2018, 126 (2), 797–849.

- Lise, Jeremy and Jean-Marc Robin**, “The macrodynamics of sorting between workers and firms,” *American Economic Review*, 2017, 107 (4), 1104–35.
- Loecker, Jan De and Jozef Konings**, “Job reallocation and productivity growth in a post-socialist economy: Evidence from Slovenian manufacturing,” *European Journal of Political Economy*, 2006, 22 (2), 388–408.
- Melitz, Marc J and Sašo Polanec**, “Dynamic Olley-Pakes productivity decomposition with entry and exit,” *The Rand Journal of Economics*, 2015, 46 (2), 362–375.
- Menzio, Guido, Irina A Telyukova, and Ludo Visschers**, “Directed search over the life cycle,” *Review of Economic Dynamics*, 2016, 19, 38–62.
- Moscarini, Giuseppe and Fabien Postel-Vinay**, “The cyclical job ladder,” *Annual Review of Economics*, 2018, 10, 165–188.
- Restuccia, Diego and Richard Rogerson**, “Misallocation and productivity,” *Review of Economic Dynamics*, 2013, 16 (1), 1–10.
- Sorkin, Isaac**, “Ranking Firms Using Revealed Preference,” *Quarterly Journal of Economics*, 2018, 133 (3), 1331–1393.
- Stoyanov, Andrey and Nikolay Zubanov**, “Productivity spillovers across firms through worker mobility,” *American Economic Journal: Applied Economics*, 2012, 4 (2), 168–98.
- Syverson, Chad**, “What determines productivity?,” *Journal of Economic literature*, 2011, 49 (2), 326–65.
- Topel, Robert H and Michael P Ward**, “Job mobility and the careers of young men,” *The Quarterly Journal of Economics*, 1992, 107 (2), 439–479.