

NBER WORKING PAPER SERIES

WORKER FLOWS, ENTRY AND PRODUCTIVITY IN THE NEW ZEALAND CONSTRUCTION
INDUSTRY

Nathan Chappell
Adam B. Jaffe
Trinh Le

Working Paper 24376
<http://www.nber.org/papers/w24376>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2018

This paper is based on Motu Working Paper Nos. 16-08 and 18-02. We thank Richard Fabling, Dean Hyslop, Dave Maré and Lynda Sanderson for valuable discussion and comments. We also thank seminar participants at Motu Research, at the 2017 New Zealand Association of Economists conference, and at a Productivity Hub roundtable event. Funding was provided by BRANZ. All opinions and errors are ours alone. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Nathan Chappell, Adam B. Jaffe, and Trinh Le. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Worker Flows, Entry and Productivity in the New Zealand Construction Industry
Nathan Chappell, Adam B. Jaffe, and Trinh Le
NBER Working Paper No. 24376
March 2018
JEL No. D24,J63,L74

ABSTRACT

The 21st century global decline in productivity growth is not well understood. One possible contributor is a decline in economic dynamism. We explore the contribution of firm formation and employee movement to productivity using administrative data on the population of New Zealand construction firms from 2001-2012, along with linked data on their employees and working proprietors, to study the relationships among entry, worker flows and firm productivity. Entrants are more productive than pre-existing firms. Firms that enter and stay exhibit a persistent productivity advantage that averages about seven percent, but which grows as experience accumulates. We find that job churn is prevalent in construction, with around 60 percent of firm-worker pairs not existing previously or not existing subsequently. Firms with new employees are more productive than those with no change in workforce, in part because of knowledge flows from other construction firms. In our preferred specification, with firm fixed effects, a standard deviation increase in the productivity of new employees' previous firms is associated with a 0.6 percent increase in productivity. The entry and worker-knowledge-flow phenomena are distinct, in that the entry effect is not explained by employee composition, and non-entrant firms also benefit from worker knowledge flows.

Nathan Chappell
Motu Economic and Public Policy Research
Level 1, 97 Cuba St.
Wellington 6011
New Zealand
Nathan.Chappell@motu.org.nz

Trinh Le
Motu Economic and Public Policy Research
PO Box 24390
Wellington 6142
New Zealand
trinh.le@motu.org.nz

Adam B. Jaffe
Motu Economic and Public Policy Research
PO Box 24390
Wellington 6142
New Zealand
and Queensland University of Technology
and also NBER
adam.jaffe@motu.org.nz

1 Introduction

From Josef Schumpeter in 1942 (Schumpeter, reissued 2013) to the most recent OECD report on the NZ economy (OECD 2017), economists have emphasized the dynamic reallocation of workers and other resources, including the birth and death of firms, as key determinants of aggregate productivity growth. The importance of dynamism to productivity has gained added urgency with the global slowdown in productivity growth over the last 15 years, with some suggesting a decline in new firm formation as a contributor to declining productivity performance (e.g. Haltiwanger, et al, 2015). Because productivity is, by definition, a residual, it is hard to say what attributes of new firms or new workers are the source of higher productivity. In this paper, we use unusually detailed micro data on firms and workers to see if productivity performance can be traced to the history of workers and firms. We focus on the construction sector because of its importance, its historically poor productivity performance, and the prevalence of small firms in which we expect to have the best chance of finding productivity effects of individual employees. We find that much of the superior performance of entrants cannot be explained even with our rich data--new firms have higher productivity and more rapid productivity growth, and this difference remains after controlling for observables. But we do find that employees' previous history does affect firms' productivity, for both new entrants and continuing firms. In particular, employees previously employed at high-productivity construction firms 'bring with them' a portion of the productivity advantage of previous employers.

The finding that new firms are more productive than exiting firms is common in the international literature (Aw et al. 2002; Foster et al. 2008) and is generally seen as a manifestation of Schumpeter's 'creative destruction'; underperforming firms are unprofitable and so leave in a process of natural selection (Jovanovic, 1982). One characteristic of new firms is that their employees have to have come from somewhere else. We expect that a firm's productivity is, in part, a result of its employees' human capital, so it is natural to wonder if the productivity of entrants is related to their employees. Further, if the higher productivity of entrants is indeed related to the employees that they attract, it is then natural to ask if ongoing firms could get the same productivity advantage if they could attract the same kind of employees. In this paper we explore empirically the interaction of the entry and exit of firms with the movement of employees across firms, both to entrants and between existing firms. We investigate this using linked employer-employee data on all construction-industry jobs in New Zealand from 2001 to 2012. We hypothesize that the flows of workers across construction firms is a source of knowledge transfer, and that these knowledge flows are part of the reason new construction firms perform well. The construction industry is particularly suitable for this analysis, because the small size of most firms likely increases the importance of individual workers with respect to firm performance.

Our study is related to two main strands of literature. The first strand documents the extent of firm churn and its correlation with aggregate productivity growth. The international evidence suggests that churn's contribution tends to be small but positive. This holds for manufacturing in the United States (Bartelsman and Dhrymes, 1998), manufacturing and mining in Israel (Griliches and Regev, 1995), manufacturing in Chile (Liu, 1993), and manufacturing in Taiwan (Aw et al., 2001). Broad overviews are given by Bartelsman et al. (2004) and Kocsis et al. (2009), who document the importance of firm churn in reallocating resources in developed and developing countries, and its positive correlation with productivity growth. In New Zealand, Doan et al. (2012) document the prevalence of firm churn and find that it is procyclical, while Law and McLellan (2005) find that net firm entry is negatively associated with industry labor productivity. In contrast, Maré et al. (2016) look at a different time period and consider multi-factor productivity (MFP) growth. They find that although entrants and exiters tend to be less productive than continuers, the net contribution of firm churn tends to be positive.

This literature highlights two main channels through which churn affects industry productivity. The first is the direct channel of reallocating resources towards the most productive firms. Firms that are going to exit tend to suffer the 'shadow of death' of lower productivity (Griliches and Regev, 1995), and so their exit releases the firms' capital and labor to be used by more productive entrants and incumbents. If exiting firms tend to be less productive than entrants, then firm churn will directly increase productivity. If entrant firms are also more productive than continuers then firm churn will boost productivity by an even larger amount. The second channel through which churn may improve productivity is indirect. The threat of entry creates competitive pressure for incumbent firms. Increased competition may then spur innovation and, ultimately, productivity growth. Together, this churn underlies the creative destruction that Schumpeter believed was the 'essential fact about capitalism' (Schumpeter, reissued 2013).

A second strand of related literature looks at the flows of employees across firms, and the extent to which this correlates with firm performance. Thanks to the increased availability of linked employer-employee data, research into worker flows has been expanding. Yet, few studies have linked these flows to firm performance. Bjelland et al. (2011) document these flows in the US, showing that they make up around 4 percent of quarterly employment and 29 percent of main job separations; that they are procyclical; that employees with less education tend to switch employers more; and that over half of employee switches are between industries (using the broadest grouping of industries). Other studies also show that job flows are large, with a lot of dynamism in various economies (Hopenhayn and Rogerson 1993, Davis et al. 2006). Maliranta et al. (2009) use Finnish data to examine whether workers spread R&D-generated knowledge when moving between firms. They conclude that knowledge from new employees can be readily implemented by firms, because hiring an employee previously in an R&D role boosts firm

performance only when moving the employee into a non-R&D role. Campbell et al. (2012) examine employer-employee flows in US legal services, and find that wealthier employees are less likely to leave a firm but are more likely to create a spin-out firm conditional on leaving. They also find firm performance is hurt more when an employee leaves to form a spin-out firm than to become an employee at another firm.

Our study makes a contribution to both strands of literature by exploring the relationship between worker flows and firm productivity. We also examine the extent to which the positive contribution of firm churn can be explained by worker flows, by looking at entrant effects after controlling for the past productivity of new workers. Our study brings these two important strands of literature together. It contributes to the 'firm churn' literature by providing new evidence that cannot be gained by examining firm-level data alone, and to the 'worker flows' literature by offering new insights into the pathway through which workers contribute to firm productivity, which is not possible by looking at individual-level data alone.

To preview our findings, we show that

- Entrant firms are observationally different from firms that were active in construction at the start of this period. They are more productive on average, they grow faster, and their productivity continues to improve over time. Accounting for observable differences in the composition of their work force reduces the 'entrant effect' somewhat but does not eliminate it.
- The productivity advantage of entrants is smaller in their first year, and the productivity performance of entrants then improves for at least a few years after entry.
- Workers frequently move in and out of jobs in construction, with an overall job churn rate of around 60 percent in most years. Firms (new or existing) that have new workers in a given year show higher productivity in that year. More generally, firm growth is associated with firm productivity improvement, though we cannot say anything about causality in this relationship.
- Firms whose work force previously worked at another high-productivity construction firm are themselves more productive, and self-report higher rates of innovation from their workers. This worker-embodied-productivity effect is small but not trivial and is statistically significant.

2 Data¹

This study uses rich administrative data from Statistics New Zealand's Integrated Data Infrastructure (IDI) and Longitudinal Business Database (LBD), which are collections of datasets containing longitudinal information on individuals, households and firms.² We use the work of Fabling and Maré (2015a), who process these data to estimate labor input in full-time equivalent (FTE) units, and to identify individuals who are working proprietors at a firm. This identification allows us to separately consider the movements of employees and working proprietors at construction firms. This is important because a large proportion of construction firms in New Zealand have no workers other than the proprietor(s), and working proprietors may contribute to firm productivity differently than other employees.

The IDI provides measures or allows construction of measures of worker characteristics including age, engagement in on-the-job training, engagement in post-secondary education, migration history and visa status. We supplement these observed characteristics with estimates of employees' unobserved worker quality from Maré and Hyslop (2006) and Maré, Hyslop and Fabling (2017). Maré and Hyslop estimate log wage regressions with observed characteristics, firm fixed effects, and employee fixed effects. They and we interpret the employee fixed effects from these regressions as a measure of unobserved aspects of an employee's quality as a worker.

Each firm in the LBD is assigned a predominant industry, using Australia and New Zealand Industry Classification (ANZSIC) 2006 codes. These allow us to identify construction firms, as well as worker flows between and within industries. We identify entrants and exiters by observing employment data over time in the IDI. A firm is an entrant if it had no labor input (either employees or working proprietors) in the previous year and is an exiter if it has no labor input in the following year.

To explore the productivity of firms, we calculate firms' real gross output, capital services, intermediate materials, and labor input. Real output, capital, and intermediate materials are measured by deflating dollar amounts captured in the LBD. The production dataset is not comprehensive, because of missing data and restrictions to ensure consistency and accuracy; firms with useable production data account for 62 percent of total gross output across all industries that are covered by production data. Nonetheless, the productivity coverage is substantial, and allows us to examine the productivity of firms that workers move between. We warn that the labor and capital measures are less reliable for working-proprietor firms than for employing firms (see Fabling and Maré 2015b for details).

¹ More detail on data construction is provided in the Data Appendix.

² See the disclaimer at the front of this paper for information on the conditions of access. The analysis in this paper uses the 20141205 archive of the IDI.

In addition to measuring current firm productivity, we use these data to measure knowledge flows from firm to firm. We do this by looking at the productivity of new workers' previous firms, as detailed in Section 4.2.

In part of our analysis we supplement the above administrative data with the 2005, 2007, 2009, and 2011 Business Operations Surveys (BOS). These surveys, run by Statistics New Zealand and targeted at firms with six or more employees, include questions on firms' sources of new ideas, which we use to corroborate our measures of knowledge flows.

3 Firm entry and productivity

Our framework for looking at productivity is the augmented production function, which models firms' output as a function of its inputs and various productivity 'shifters'. These regressions take the following form:

$$y_{jt} = \alpha_1 k_{jt} + \alpha_2 l_{jt} + \alpha_3 m_{jt} + \omega Z_{jt} + \beta X_{jt} + \delta_j + \gamma_t + \varepsilon_{jt} \quad (1)$$

where j denotes firm and t denotes year. The lower-case y , k , l , and m represent the log of gross output, capital, labor and intermediate inputs. Z_{jt} denotes a vector of other firm-year specific control variables including indicators for having production data from the Annual Enterprise Survey (AES, typically larger firms), and an indicator for being in Canterbury after the 2011 earthquake.³ X_{jt} represents one or more variables that are the focus of the study, hypothesized to shift the productivity of the firm. The terms γ_t and δ_j correspond to year and firm effects. We include year-specific intercepts in all regressions. We explore different approaches to unobserved firm effects, including a variety of industry dummies as well as firm fixed effects. The error term is ε_{jt} and captures variations in gross output not explained by the model. Hence our coefficient estimates of interest are in the vector β , and capture the associations of workers' past productivity with a firm's current productivity.⁴

We begin by exploring the productivity performance of entrants relative to that of continuing firms, based on the complete sample. This is an empirically complex question, because firms' entry and related decisions are likely correlated with unobserved attributes. We therefore do not attempt to build a causal model of behaviour, but rather simply to characterize the regularities in the data.

³ In September 2010 and February 2011 the Canterbury region, where about 13% of the New Zealand population resided, experienced severe earthquakes which resulted in widespread damage to buildings and infrastructure. Extensive rebuilding has since started.

⁴ To conserve space, the coefficients for the non-augmenting terms are not reported in the results tables. Appendix Table 6 provides complete regression results for columns 2 and 3 of Table 1; the unreported coefficients for other regressions are all similar.

The first column of Table 1 asks simply whether the productivity of the firms observed entering during our data period differs, on average, from the productivity of those firms that are present in the dataset in the first year. The answer is 'yes'--for given level of inputs and controlling for industry these firms produce on average about six percent more output. We will see that a large overall difference for these firms is a robust feature of these data.

The subsequent columns in Table 1 show how the performance of entrants evolves over time after entry. Column 2 includes in the regression all firms in the dataset, and allows for a productivity difference in each of the first four years of life, and the last year. The omitted group, and hence the group against which each of the dummy coefficients is calculated, is all years that are not years 1-4 and are not a final year. The results suggest that the first few years of a firm's existence are comparatively high-productivity years. But it is important to think carefully about what these coefficients actually measure. Firms that are not observed as entrants never have an early year observed in our data. So the comparison underlying the coefficients in the first column is between the early years *of entrant firms*, and a mixture of (a) years 5 and up for entrants, and (b) *all* years for non-entrants. Further, some firms enter and then exit. The firms for which year 2, year 3 or year 4 is observed are each somewhat different. To the extent that the entrants are intrinsically different from the non-entrants, and/or firms that enter and then exit differ from persistent entrants, column 2 is hard to interpret.

To unpick this puzzle, columns 3-6 of Table 1 are estimated including only those firms that are observed as entrants and then persist for at least 4 years.⁵ This means that the coefficients are not influenced by the *overall* difference in productivity between the entrant group and the non-entrant group, and do not contain the effects of firms that enter and then exit relatively quickly. Column 3 is just like column 2, but estimated on this smaller sample. We find that the positive and significant coefficient in the first year goes away, suggesting that that result in the column 2 is dominated by entrants overall being more productive than non-entrants, rather than by what goes on in entrants' first year.⁶ But the interpretation of the results in column 3 is still made somewhat murky by the fact that we observe some entrants for 4 years, some for 5, some for 6, etc. The coefficient on year 1 is the average difference between year 1 and all years greater than 4. If firms that appear in the data for different durations differ from each other, those across-firm differences will be affecting the results in column 2 in an unknown way.

Column 4 explores this dynamic further by estimating the model with firm fixed effects. This means that the coefficients for each dummy represent the average difference for the observations

⁵ Appendix Table 1 presents simple statistics for the variables appearing in Table 1, for this main sample of firms ever seen entering which also appear at least four times.

⁶ Other results (not reported) suggest that there may also be a tendency for firms that enter and then exit quickly to show relatively high productivity while they last. This is likely a mixture of data problems and specialized firms that are created for a single project and then dissolved. Because it is not our main interest, and these firms by their nature cannot have a lasting impact on industry productivity, we focus on those firms that remain active at least 4 years.

in that group between that observation and the mean for the corresponding firm. The coefficient on the first year is now negative and statistically significant. This means that for firms that enter at some point in the observed period, the first year we observe them exhibits about seven percent lower productivity than the average in the data for that firm. This below-average performance falls to about one percent in years 2-4.

In Section 4.2 we explore the relationship between firm productivity and the previous work experience of employees. Without going into the details of that formulation, the last two columns simply test whether inclusion of those effects changes the results regarding entrant productivity dynamics. For the most part, they do not. That is, column 5 is similar to column 3 and column 6 is very similar to column 4. Thus the dynamics of entry are largely distinct empirically from the effects of workforce composition. Finally, we note that the final year a firm is observed shows productivity consistently about 13 percent below non-exit years, and this estimate is highly robust to the different specifications.

The fixed effect for each firm (its overall mean) is estimated, of course, on the data observed. For some firms this is 5 years; for some it is more. For a firm that is in the data for 5 years, we do not know what its firm mean would be if it were around for 10. Ultimately, since we have unbalanced data and we do not know why firms come and go, there are limits to how much we can say about the underlying dynamics. But we can note two strong tendencies, even if we cannot explain them: firms that are observed to enter sometime in the period have persistently higher productivity than non-entrants. And the productivity observed in a firm's first year, for entrants that continue to produce for at least 4 years, is lower than the average productivity eventually observed for the firm. This suggests that the entrants are intrinsically different firms than the non-entrants, and that their superior performance is not an artefact of the start-up situation. Indeed, the data are consistent with a shake-down or learning period over which entrants gradually achieve their potential.

4 Firm productivity and experience of new workers

4.1 Importance of new labor in construction

Before turning to the relationship between productivity and workforce composition, we pause to simply describe the extent of new labor at construction firms. Figure 1 shows the job entry, exit, and churn rates over 2001-2013 for workers in construction, whether they are employees or working proprietors. For this purpose, a 'job' is a match between a specific worker and a specific firm: a new (dying) job is defined as the first (last) year in which such a match is observed in our data. The job entry rate is the proportion of all jobs that are new jobs, the job exit rate is the proportion of jobs that are dying jobs, and the job churn rate is the sum of the job entry

rate and the job exit rate. The job churn rate is large; in most years about 60 percent of jobs either never existed in the past, or will not exist again in the future. This drops to just under 50 percent in the wake of the global financial crisis, consistent with the general finding that job churn is procyclical (Lazear and Spletzer 2012, Bjelland et al. 2011).⁷

Figure 1 counts all jobs equally regardless of the fraction of full-time, and new jobs are for less time than others on average. Figure 2 shows that the proportion of labor in construction that is new on an FTE basis varies from 19 percent (in 2005) to around 10 percent (2010). The figure also shows where new labor comes from: other construction firms, firms in other industries, or not in the job data in the previous year (the left out category, not graphed).⁸ In most years, around five percent of all labor comes from new jobs where the worker was in construction the previous year. A similar rate holds for labor from other industries.

Figure 3 shows the 4-year history and future of all people working in construction between 2005 and 2010. This further highlights the dynamism of the labor market in construction; among those employed, fewer than 40 percent of people held the same job four years prior and only around 40 percent held the same job four years after. Interestingly, the second most common category is the miscellaneous omitted group. This means that many current workers in construction have pasts and futures that involve some combination of education, overseas travel, several part-time jobs with less than 0.25 FTE worked at each one, and other activities.

Panel B replicates the above for new workers at entrant firms. A new worker cannot have worked at the same entrant firm in previous years, and so the remaining three categories have higher proportions mechanically. However, note that entrant workers are especially likely to come from other construction firms; in the pre-period the proportion coming from construction diverges from the proportion coming from another industry, while the pattern is the opposite in Panel A. This pattern suggests one reason entrant firms may be more productive – they source their labor from other construction firms. We explore this hypothesis statistically in the next subsection.

4.2 Firm productivity and experience of new workers

We turn next to a set of augmented production function regressions that explore the association between firm productivity and the introduction of new employees with varying experience in terms of the productivity of their previous employers. The variables capturing new workers' average MFP at previous firms is split into employees and working proprietors, and is a

⁷ The pattern is similar when defining a new job as a relationship that did not exist the previous year, and an exiting job as one that does not exist the following year.

⁸ A person's labor 'from construction' is calculated by weighting her current labor input by her labor input at construction firms the previous year. For example, if someone works 0.5 FTE this year, and worked 0.5 FTE in construction the previous year, then she contributes 0.25 FTE from construction. Labor from other industries is calculated analogously.

double-FTE weighted measure⁹: first, we weight previous MFP by the FTE worked at that firm, because a worker will transfer less knowledge when only working at a previous firm for a short time. Secondly, we weight these past-productivity measures by the current FTE worked with the firm, because a worker will transfer less knowledge and have less impact when only working at a new job for less time. The past productivity of new workers is measured over the previous two years to reduce noise while still reflecting relatively recent experience. We scale the employee-past-productivity-experience variables by firm size, expecting that the impact of any one worker is less the larger the firm.¹⁰

In addition to these productivity-transfer variables, we include in the regression a set of dummy variables that simply capture whether the firm in that year had any new workers from various sources. We include these because firms that are adding new employees in any given year are likely to be experiencing other environmental factors that may be correlated with productivity. We interpret the coefficients on these dummy variables as capturing such incidental dynamic factors, so that the coefficients on the previous MFP experience of the workers can be interpreted as associated with those workers rather than background conditions.

The first column of Table 2 reproduces, for comparative purposes, the basic entrant-effect regression of Table 1.¹¹ The second column shows the coefficients on the workforce composition variables excluding the entrant effect, and the third column includes both effects. Comparing the three columns demonstrates some basic results. First, the entrant effect and the knowledge transfer effect are both present, and they do not interact strongly. That is, the magnitude of each effect is not greatly changed by inclusion of the other effect. Second, there is a clear effect of bringing in workers from productive firms, and the knowledge-transfer interpretation of this effect is strengthened by the fact that the productivity benefit is greater when the previous high-MFP employer was a construction firm. Indeed, for employees the effect of productivity of previous employers outside of the construction industry is not statistically significant.

Because the knowledge-flow variables are weighted by the employees' FTE in both the productivity-sending firm and the productivity-receiving firm, they can be interpreted as a kind of fractional effectiveness of transfer. That is, the coefficient of 0.32 in the third column, for example, means that if we consider two construction firms with differing productivity levels, and then imagine hypothetically producing two new firms each of which is constructed by simply

⁹ The 'FTE' measure for working proprietors is not an accurate labor-input measure. Anyone receiving self-employed income is given a labor-input measure of 1, and this is scaled by the number of working-proprietor jobs in the case of several self-employment relationships. We also emphasize again that the employee FTE measures tend to overestimate actual labor input, as discussed in the data appendix. See Fabling and Maré (2015a) for details.

¹⁰ Note that past MFP is measured by the Cobb-Douglas production function residuals of Fabling and Maré (2015b) and are estimated separately for industries roughly corresponding to Statistics New Zealand's ANZSIC 2006 level 2 industries. For construction, there are three level-2 industries: building construction, heavy and civil engineering construction, and construction services.

¹¹ Appendix Table 2 presents simple statistics for the variables appearing in Table 2, for this main sample of firms which appear at least four times.

transferring all of the employees from one of the original two firms, the difference in productivity between the two new firms would be 38 percent ($=\exp(0.32)-1$) of the difference between the original firms. Of course, in reality, firms will be a mixture of old and new employees. If only half of the employees at the new firms come from other firms, then the effect would be half as large.

The estimate for working proprietor past productivity in construction is also statistically significant. For previous experience in the construction industry, the magnitude is markedly smaller for working proprietors than for employees; for experience outside of construction it is larger for working proprietors and is statistically significant. This may reflect the likelihood that employees coming from outside construction may lack relevant skills, while a proprietor of a new construction firm likely has significant construction skills even if their previous employment was not in construction.

Finally, the regression also includes dummy variables for a firm's having any new workers from a given source.¹² The positive estimates suggest that firms with new employees are more productive than firms with a static work force. Note that this effect is diminished (but still present) when the entrant dummy is included in the regression. This makes sense, since entrants by definition have new employees. We interpret this effect as likely reflecting that when unobserved good things happen to a firm, it is likely to grow and it is likely to experience a positive productivity shock. Hence we cannot say anything about causality in this relationship. We believe that the importance of these variables lies mainly in capturing the effect of unobserved good news on both productivity and employment growth, so that the effect associated with employment of people from high-productivity firms are less likely to be spurious. This issue is discussed further in Section 5.

The results in the third column demonstrate an association between the employment of workers taken from high-productivity firms and high productivity of the newly-employing firm. We would like to interpret this effect in terms of the employees bringing with them knowledge about high-productivity practices. But another interpretation is that the firms that are able to attract new employees away from high-productivity firms are simply better firms, and so their higher productivity reflects that difference in the innate quality of the firm, not any effect brought by the new workers. Appendix Table 3 examines the extent to which above-average workers tend to work at above-average firms. Using the worker and firm fixed effects estimates of Maré, Hyslop and Fabling (2017), it groups firms by whether they are 'good' firms (high-paying regardless of who their employees are) and shows the proportion of their workers who are 'good' (highly paid wherever they work). The table shows that 'good' firms are somewhat more likely to have 'good' workers, both in terms of employees overall and in terms of new hires, though the differences are not large--just a few percentage points in each case. That is, the firm characteristics represented

¹² Above a threshold of 0.15 FTE worked when looking at new employees.

by the unobserved δ_j in Equation (1) may be correlated with the variables capturing worker knowledge flows; in that case our inability to include δ_j in the regression biases our estimate of the knowledge flow effects.

We explore the impact of this 'sorting' phenomenon on our results first by using different econometric approaches to handling the effect of δ_j . Column 4 is estimated using firm fixed effects. This is an appropriate estimation strategy if δ_j is constant across t for given j .¹³ As expected, the result of this estimation approach is to diminish the estimated magnitude of the knowledge transfer effect, but it is still present for both employees and working proprietors coming from previously high-MFP firms. The estimate now suggests that about nine percent of the difference in productivity between source firms is reflected in the productivity of employee-acquiring firms. Another way to interpret the magnitude is that a standard deviation increase in the productivity of new employees' previous firms is associated with a 0.6 percent increase in productivity (.064 * .088 = .006).

Puzzlingly, there is now a negative and statistically significant effect of hiring workers from high-productivity firms outside of construction. While it is easy to see why knowledge from non-construction firms might not be useful, it is hard to see why it would actually be destructive of productivity. We suspect that this effect may reflect the productivity consequences of hiring workers with particular industrial backgrounds.

The last column takes as its dependent variable the change in the log of output rather than the log itself. This approach causes δ_j to drop out of the equation, eliminating the bias but also discarding a lot of information so that it becomes harder to get precise estimates. This loss of precision is seen in the fact that most of the estimates in the last column are not statistically significant. Most are, however, qualitatively similar to the fixed-effect results. One interesting result from the first-difference regression is that the coefficient on 'ever an entrant' is positive and statistically significant. Since the dependent variable in this regression is the *change* in log output, the coefficient of .015 tells us that these firms see productivity *growth* each year that is 1.5 percentage points higher than the other firms. This reinforces the suggestion that these firms are fundamentally different than the pre-existing firms.

A different way to control for whatever amount of employer-employee sorting may be occurring is to add augmenting variables to the production function that capture employee characteristics. Column 1 of Appendix Table 4 repeats our main specification but controls for the average unobserved and observed skills of a firm's new workers. Unobserved skills are measured with estimated worker fixed effects, while observed skills come from gender-year specific age

¹³ The 'ever an entrant' variable is dropped from this regression because it is fixed over time for a given firm.

profiles.¹⁴ The average worker skill variable is positive and statistically significant for new employees, and positive though mostly insignificant for new working proprietors, showing that firms which hire more skilled workers tend to be more productive. This addition does not, however, materially change the estimates on the four knowledge-flow variables. Column 2 includes firm fixed effects in our augmented production function. The results are similar to column 4 of Table 2, supporting the knowledge-flow interpretation of the MFP-transfer results.

For robustness, columns 3 and 4 of Appendix Table 4 replicate columns 3 and 4 of Table 2 but weight each firm by its gross output. This approach gives more importance to larger firms, allowing the interpretation of the result for the average unit of output in construction rather than for the average firm. The results are somewhat smaller for the ‘past productivity from construction’ variables, again suggesting that knowledge flows are less important for larger firms. This holds even after scaling key covariates by firm size.¹⁵

While the results so far allow entrants to have a different productivity level than other firms, they constrain the MFP-transfer effects to be the same in the two groups. The last column of Appendix Table 4 explores whether knowledge flows matter more for new firms, by limiting the sample to entrants in construction. The regression excludes firm fixed effects because few firms enter more than once.¹⁶ The estimates for three of the past-productivity variables are slightly larger than the corresponding estimates for all firms in column 3 of Table 2, while the estimate for ‘new employees coming from non-construction firms’ is twice as large. Overall, the results suggest the phenomenon of knowledge flows is similar in entrant firms.

Table 3 replicates the main results (columns 3 and 4 of Table 2) separately for the main sub-industries of construction. Note first that in every sub-industry, there is a reasonably large and statistically significant positive effect for the entrant firms; for ‘heavy and civil construction’ this effect is roughly twice as large as for the industry overall. The MFP-transfer results are similar for ‘building construction’ and ‘construction services’, but are noisy and imprecisely estimated for ‘heavy and civil construction’ due to the small number of firms. The ‘past-productivity from construction’ estimates for ‘construction services’ are statistically significant and larger in magnitude than the pooled-industry estimates, suggesting knowledge flows are especially important in this sub-industry. Construction services firms tend to be especially small, which amplifies the importance of worker flows to the firm’s performance. The patterns are similar for

¹⁴ See the data section for more details. We use one skills variable by adding together each person’s work fixed effect and each person’s predicted benefit to wages because of the gender-year specific age profile. We then standardize this measure before including it in our regression.

¹⁵ We also ran regressions separately for small firms (under the 75th percentile of firm size, which is about 2.2 FTE) and large firms (at and above the 75th percentile of firm size). Our main estimates are positive and statistically significant for small firms, and smaller and insignificant, though still positive, for large firms.

¹⁶ Although this is possible, because we define entrants as firms with no labor input in the previous year.

'building construction', though smaller in magnitude and statistically insignificant in the specification with firm fixed effects.¹⁷

Finally, Appendix Table 5 replicates columns 3 and 4 of Table 2, for all of construction and the three sub-industries, but uses a different measure of knowledge flows. These new measures consider separately labor from productive and unproductive firms, for construction and other industries. The main patterns are similar, where new employees from productive construction firms matter more than those from productive non-construction firms (though this does not hold for new working proprietors).¹⁸

4.3 Employee movement and new ideas

We believe that the regression models in the previous section are strongly supportive of the hypothesis that employees from high-MFP firms can to some degree transfer the knowledge underlying that high MFP to other firms when they change jobs. But of course it is hard to know for sure because we do not observe the knowledge transfer itself. To shed some additional light on the issue, Table 4 explores whether our measure of employees' past productivity is associated with firms' reporting new ideas. The dependent variables come from the innovation modules of the 2005-2011 BOS, and so the regressions are limited to the much smaller sample of construction firms answering these questions. The first column shows that firms with higher employee past productivity are more likely to report new ideas for innovation from new staff: consider a firm that gains new employees making up half its labor force, where the employees are coming from construction firms ten percent more productive than the industry average. The coefficient of 2.261 means that such a firm is 11 percentage points more likely to report new ideas from new staff, on average and holding all else constant ($2.261 * 0.5 * 0.1 = 0.113$).

A similarly strong relationship is found when looking at new ideas from firms in the same industry, while estimates are smaller and statistically insignificant when looking at reporting of ideas from old staff and ideas from firms in other industries. This is reassuring, and gives further evidence that our measures of past productivity are capturing knowledge flows from other firms. When looking at the past productivity of non-construction workers, the coefficient estimates are smaller and statistically insignificant. This suggests that within-industry knowledge flows are

¹⁷ We also ran regressions estimating the 1-year change in log gross output based on the change in the log of inputs and the level of the past-productivity variables and their associated dummy variables. Estimates, not reported, are similar to our firm-fixed-effects specifications, though slightly smaller in magnitude. The ever-entrant dummy is positive, suggesting much older firms also grow more slowly (1.4 percent less per year for all of construction).

¹⁸ We also added variables capturing the average age of new employees and working proprietors, and the proportion of a firm's labor force with industry training. The unreported results are similar, showing that our estimates are not driven by the age of new workers or their industry training.

especially important, and that part of the mechanism is through innovative ideas and not merely productivity embodied in a worker that will be lost when the worker moves on.¹⁹

One caveat is that the BOS survey only targets firms with six or more employees. More than 90 percent of construction firms do not meet this threshold. The last section suggested that the transfer of productivity is stronger in small firms; it is not implausible that the innovation effects would be similarly as strong or stronger in smaller firms.

5 Summary and discussion

This paper has analysed how the flows of firms into markets, and workers into firms, interacts with the productivity of firms. Although international literature suggests firm churn is a sign of a healthy economy, the potential for knowledge flows through workers is especially high in New Zealand's construction industry due to the large number of very small firms.

We find that firms observed to enter during this period are systematically and persistently more productive than previously existing firms. Looking at the time path of productivity of entrants, we find them to be modestly less productive in their first year than they later become. We believe that these patterns suggest that entrant firms have intrinsic advantages over previously existing firms, but it does take them a year or more to build their productivity to its true potential.

We have shown that workers regularly change jobs in construction: in most years between 2001 and 2013 around 60 percent of worker-firm pairs either did not previously exist or will not exist the following year. In terms of overall averages, this movement of workers is associated with higher productivity. Firms that gain or lose any labor are more productive than static firms, and firms that gain labor from other construction firms are especially productive.

By looking at the past productivity of new workers' previous firms, we also find evidence that workers transfer productivity when changing firms; the average firm is more productive when more of its labor is made up of new workers who were previously employed at more productive firms. Comparing the results with and without firm fixed effects suggests that much of the impact is a matching effect--high quality firms attract and employ workers from other high-quality firms. There is, however, a portion of the impact that remains after controlling for this effect, and workers from construction seem to transfer productivity to a much greater extent than workers from other industries. These results do suggest that there is a small but statistically significant effect that can be associated with knowledge flowing from firm to firm through worker flows.

¹⁹ Though note that these patterns do not hold when measuring past productivity as in Appendix Table 5. One interpretation is that our main measure is better, because it correlates with these BOS variables in the expected ways. But it is also important to maintain some skepticism and remember that our measures are not perfect.

The knowledge-flow interpretation is corroborated by a finding that firms with new employees sourced from high-MFP construction firms are more likely to report getting ideas for innovation from new staff and from other firms in construction, while they are not more likely to report getting innovative ideas from old staff or from outside construction. Further, having new employees from high-MFP firms *outside* construction is not associated with reporting getting new ideas.

The meaning of this association is clouded by the fact that firms' and employees' decisions may be correlated with unobserved aspects of the environment that are correlated with productivity. We have attempted to control for this by including dummy variables for the arrival of new workers from various sources, and by exploring the relationship with fixed effects, in terms of productivity growth, and also in terms of reported innovation sourced from employees. We believe that the persistence of the effect running through these different approaches suggests there is likely something real.

Together, our results confirm that inflows of new firms and movements of employees are important to firm productivity. New firms are more productive than existing firms--briefly and perhaps meaninglessly for those that quickly disappear, but persistently and increasingly for those that do persist. This is in part because entrants enjoy the benefits of knowledge flows from other firms carried by their new employees, but the new employee/knowledge-flow effect and the entrant effect are mostly distinct effects--entrants are more productive than can be explained by the previous experience of their new employees, and non-entrants also benefit from knowledge associated with new employees.

Further research could explore whether these relationships hold and are as strong in other industries. Further work could also look at the less productive firms whose existence predates the data period. The implications for policy depend on whether old firms have obsolete technology, have market power and less incentive to innovate, or have simply grown complacent. Finally, it would be interesting to investigate whether these patterns have persisted in more recent years when the industry has been growing rapidly. It is possible that entrants and new employees who are drawn into the industry when demand surges are different in nature from entrants and new employees during more stable periods.

References

- Aw, B. Y., Chen, X., & Roberts, M. J. (2001). Firm-level evidence on productivity differentials and turnover in Taiwanese manufacturing. *Journal of Development Economics*, 66(1), 51–86.
- Bartelsman, E. J., & Dhrymes, P. J. (1998). Productivity Dynamics: U.S. Manufacturing Plants, 1972–1986. *Journal of Productivity Analysis*, 9(1), 5–34.
- Bartelsman, E. J., Haltiwanger, J., & Scarpetta, S. (2004). *Microeconomic Evidence of Creative Destruction in Industrial and Developing Countries*. The World Bank, Policy Research Working Paper Series No. 3464.

- Bjelland, M., Fallick, B., Haltiwanger, J., & McEntarfer, E. (2011). Employer-to-Employer Flows in the United States: Estimates Using Linked Employer-Employee Data. *Journal of Business & Economic Statistics*, 29(4), 493–505.
- Campbell, B. A., Ganco, M., Franco, A. M., & Agarwal, R. (2012). Who leaves, where to, and why worry? employee mobility, entrepreneurship and effects on source firm performance. *Strategic Management Journal*, 33(1), 65–87.
- Davis, S. J., Faberman, R. J., & Haltiwanger, J. (2006). The Flow Approach to Labor Markets: New Data Sources and Micro–Macro Links. *The Journal of Economic Perspectives*, 20(3), 3–26.
- Doan, T., Devine, H., Nunns, P., & Stevens, P. (2012). Firm entry and exit in New Zealand industries. Wellington: Ministry of Economic Development.
- Fabling, R. (2011). Keeping it Together: Tracking Firms on New Zealand’s Longitudinal Business Database. Motu Economic and Public Policy Research Working Papers 11_01.
- Fabling, R., & Maré, D. C. (2015a). Addressing the Absence of Hours Information in Linked Employer-Employee Data. Motu Economic and Public Policy Research Working Papers 15_17.
- Fabling, R., & Maré, D. C. (2015b). Production Function Estimation Using New Zealand’s Longitudinal Business Database. Motu Economic and Public Policy Research Working Papers 15_15.
- Foster, L., Haltiwanger, J., & Syverson, C. (2008). Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *The American Economic Review*, 98(1), 394–425.
- Griliches, Z., & Regev, H. (1995). Firm productivity in Israeli industry 1979–1988. *Journal of Econometrics*, 65(1), 175–203.
- Haltiwanger, John, Ryan Decker, and Ron Jarmin (2015) "Top ten signs of declining business dynamism and entrepreneurship in the US." Kauffman Foundation New Entrepreneurial Growth Conference.
- Hopenhayn, H., & Rogerson, R. (1993). Job Turnover and Policy Evaluation: A General Equilibrium Analysis. *Journal of Political Economy*, 101(5), 915–938.
- Jovanovic, B. (1982). Selection and the Evolution of Industry. *Econometrica*, 50(3), 649–670.
- Kocsis, V., Shestalova, V., Wiel, H. van der, Zubanov, N., Lukach, R., & Minne, B. (2009). Relation entry, exit and productivity: an overview of recent theoretical and empirical literature (CPB Document No. 180). CPB Netherlands Bureau for Economic Policy Analysis.
- Law, D., & McLellan, N. (2005). The contributions from firm entry, exit and continuation to labour productivity growth in New Zealand. New Zealand Treasury Working Paper 05/01.
- Lazear, E. P., & Spletzer, J. R. (2012). Hiring, Churn, and the Business Cycle. *The American Economic Review*, 102(3), 575–579.
- Liu, L. (1993). Entry-exit, learning, and productivity change Evidence from Chile. *Journal of Development Economics*, 42(2), 217–242.
- Maliranta, M., & Määtänen, N. (2013). Allocation and industry productivity: Accounting for firm turnover (ETLA Working Papers No. 11). The Research Institute of the Finnish Economy.
- Maliranta, M., Mohnen, P., & Rouvinen, P. (2009). Is inter-firm labor mobility a channel of knowledge spillovers? Evidence from a linked employer–employee panel. *Industrial and Corporate Change*, 18(6), 1161–1191.
- Maré, D. C., & Hyslop, D. (2006). Worker-Firm Heterogeneity and Matching: An Analysis Using Worker and Firm Fixed Effects Estimated from LEED. Wellington: Statistics New Zealand.
- Maré, D. C., Hyslop, D. R., & Fabling, R. (2017). Firm productivity growth and skill. *New Zealand Economic Papers*, 51(3), 302–326.
- OECD (2017). OECD Economic Survey of New Zealand: Boosting productivity and adapting to the changing labour market. Available at www.oecd.org/economy/surveys/economic-survey-new-zealand.htm
- Schumpeter, J. A. (2013). *Capitalism, socialism and democracy*. Routledge.

Table 1: Firms' productivity over time

	Log gross output (1)	Log gross output (2)	Log gross output (3)	Log gross output (4)	Log gross output (5)	Log gross output (6)
Ever seen entering (2001 onwards)	0.061*** (0.004)					
Firm's 1st year		0.037*** (0.005)	-0.002 (0.006)	-0.070*** (0.007)	-0.035*** (0.006)	-0.085*** (0.008)
Firm's 2nd year		0.061*** (0.004)	0.029*** (0.005)	-0.018*** (0.006)	0.023*** (0.005)	-0.020*** (0.006)
Firm's 3rd year		0.048*** (0.004)	0.017*** (0.004)	-0.019*** (0.005)	0.016*** (0.004)	-0.018*** (0.005)
Firm's 4th year		0.049*** (0.004)	0.017*** (0.004)	-0.011*** (0.004)	0.018*** (0.004)	-0.010** (0.004)
Firm's final year		-0.155*** (0.006)	-0.145*** (0.008)	-0.129*** (0.007)	-0.145*** (0.008)	-0.131*** (0.007)
Sample	Firms which appear 4+ times	Firms which appear 4+ times	Ever-entrants which appear 4+ times	Ever-entrants which appear 4+ times	Ever-entrants which appear 4+ times	Ever-entrants which appear 4+ times
Controls for new workers' past productivity	No	No	No	No	Yes	Yes
<i>Observations</i>	281,379	281,379	156,159	156,159	156,159	156,159
<i>Number of firms</i>	59,334	59,334	37,929	37,929	37,929	37,929
<i>R-squared</i>	0.885	0.886	0.853	0.665	0.854	0.665
<i>Firm FE</i>	No	No	No	Yes	No	Yes
<i>Industries included</i>	All constr. with lvl 2 FE	All constr. with lvl 2 FE	All constr. with lvl 2 FE	All constr.	All constr. with lvl 2 FE	All constr.

Notes: All counts reported in this paper have been randomly rounded to base 3 to protect confidentiality. All regressions also control for the log of capital, log of labor, log of intermediate consumption, year and region fixed effects, and include dummy variables for having production data from AES survey and for being in Canterbury in the post-earthquake period. The sample is limited to firms sized at least 0.5, and is limited to the years 2003-2012. Standard errors, in parentheses, are robust and clustered at the firm level. Simple statistics for all regression variables are presented in Appendix Table 1. Asterisks denote: *** p<0.01, ** p<0.05, * p<0.10.

Table 2: Knowledge flows and productivity

	Log gross output (1)	Log gross output (2)	Log gross output (3)	Log gross output (4)	2-yr change in log gross output (5)
Ever seen entering (2001 onwards)	0.061*** (0.004)		0.055*** (0.004)		0.015*** (0.002)
New employee mfp at prev constr firms/firm size		0.326*** (0.052)	0.320*** (0.052)	0.088** (0.035)	0.112* (0.058)
New employee mfp at prev non-constr firms/firm size		0.020 (0.024)	0.020 (0.024)	-0.039** (0.019)	-0.042 (0.053)
New WP mfp as WP at prev constr firms/firm size		0.128*** (0.014)	0.128*** (0.014)	0.030*** (0.012)	0.096* (0.058)
New WP mfp as WP at prev non-constr firms/firm size		0.074*** (0.025)	0.074*** (0.025)	0.001 (0.014)	0.103*** (0.035)
Some new employees from constr		0.022*** (0.003)	0.015*** (0.003)	0.010*** (0.002)	0.015*** (0.002)
Some new employees from other industries		0.022*** (0.003)	0.015*** (0.003)	0.010*** (0.002)	0.023*** (0.002)
Some new WPs from construction		0.112*** (0.007)	0.090*** (0.007)	0.010* (0.006)	0.082 (0.059)
Some new WPs from other industries		0.118*** (0.012)	0.098*** (0.012)	0.020** (0.010)	0.063** (0.019)
<i>Observations</i>	281,379	281,379	281,379	281,379	195,585
<i>Number of firms</i>	59,334	59,334	59,334	59,334	37,155
<i>Std. deviation of new employees' mfp at construction</i>		0.064	0.064	0.064	0.039
<i>Std. deviation of new WP mfp at construction firms</i>		0.756	0.756	0.756	0.351
<i>R-squared</i>	0.885	0.885	0.886	0.663	0.655
<i>Firm FE</i>	No	No	No	Yes	No
<i>Industries included</i>	All constr with lvl 2 FE	All constr with lvl 2 FE	All constr with lvl 2 FE	All constr	All constr with lvl 2 FE

Notes: All regressions also control for log of capital, labor, and intermediate consumption (levels in columns 1-4 and 2-year changes in column 5), year and region fixed effects, and include dummy variables for having production data from AES survey and for being in Canterbury in the post-earthquake period. The sample is limited to firms sized at least 0.5 and appearing at least four times over the period 2003-2012, and is limited to the years 2003-2012. Standard errors, in parentheses, are robust and clustered at the firm level. Simple statistics for all regression variables are presented in Appendix Table 2. Asterisks denote: *** p<0.01, ** p<0.05, * p<0.10.

Table 3: Knowledge flows and productivity, by sub-industry

	Log gross output (1)	Log gross output (2)	Log gross output (3)	Log gross output (4)	Log gross output (5)	Log gross output (6)
New employee mfp at prev constr firms/firm size	0.239*** (0.064)	0.047 (0.051)	-0.014 (0.148)	-0.107 (0.145)	0.418*** (0.081)	0.138*** (0.045)
New employee mfp at prev non-constr firms/firm size	0.027 (0.038)	-0.042 (0.029)	-0.036 (0.162)	-0.123 (0.164)	0.020 (0.031)	-0.037 (0.025)
New WP mfp as WP at prev constr firms/firm size	0.099*** (0.021)	0.023 (0.014)	0.164** (0.067)	0.081 (0.055)	0.149*** (0.018)	0.033* (0.017)
New WP mfp as WP at prev non-constr firms/firm size	0.108*** (0.033)	-0.016 (0.021)	0.052 (0.065)	0.025 (0.024)	0.049*** (0.018)	0.005 (0.017)
Some new employees from constr	0.015*** (0.005)	0.008** (0.004)	-0.016 (0.012)	0.015** (0.007)	0.018*** (0.003)	0.010*** (0.002)
Some new employees from other industries	0.024*** (0.005)	0.009** (0.004)	0.007 (0.012)	0.015 (0.009)	0.009*** (0.003)	0.009*** (0.002)
Some new WPs from construction	0.076*** (0.010)	-0.004 (0.009)	0.164*** (0.034)	0.045* (0.024)	0.091*** (0.009)	0.014* (0.008)
Some new WPs from other industries	0.087*** (0.022)	0.061*** (0.021)	0.125*** (0.048)	0.025 (0.026)	0.102*** (0.014)	0.007 (0.012)
Ever seen entering (2001 onwards)	0.077*** (0.007)		0.131*** (0.022)		0.043*** (0.005)	
<i>Observations</i>	90,690	90,690	8,577	8,577	182,112	182,112
<i>Number of firms</i>	20,025	20,025	1,683	1,683	37,626	37,626
<i>R-squared</i>	0.887	0.675	0.942	0.719	0.878	0.654
<i>Firm FE</i>	No	Yes	No	Yes	No	Yes
<i>Industries included</i>	Building constr	Building constr	Heavy & civil constr	Heavy & civil constr	Constr services	Constr services


Notes: All regressions also control for log of capital, log of labor, log of intermediate materials, year and region fixed effects, and include dummy variables for having production data from AES survey and for being in Canterbury in the post-earthquake period. The sample is limited to firms sized at least 0.5 and appearing at least four times over the period 2003-2012, and is limited to the years 2003-2012. Standard errors, in parentheses, are robust and clustered at the firm level. Asterisks denote: *** p<0.01, ** p<0.05, * p<0.10.

Table 4: Sources of new ideas and employee past-productivity

	Any new ideas for innovation from new staff	Any new ideas for innovation from old staff	Any new ideas for innovation from firms in the same industry	Any new ideas for innovation from firms in other industries
	(1)	(2)	(3)	(4)
New employee mfp at prev. constr. firms/firm size	2.261** (0.923)	1.085 (0.678)	1.957** (0.857)	0.489 (1.318)
New employee mfp at prev. non-constr. firms/firm size	-0.013 (1.017)	0.458 (1.042)	0.972 (0.849)	-0.183 (0.746)
Log firm size	0.086*** (0.019)	0.070*** (0.016)	0.027 (0.021)	0.012 (0.020)
<i>Mean of dependent variable</i>	0.637	0.797	0.413	0.213
<i>Observations</i>	477	477	459	444
<i>Number of firms</i>	201	201	198	192
<i>R-squared</i>	0.073	0.060	0.052	0.022
<i>Industries included</i>	All constr with lvl 2 FE	All constr with lvl 2 FE	All constr with lvl 2 FE	All constr with lvl 2 FE

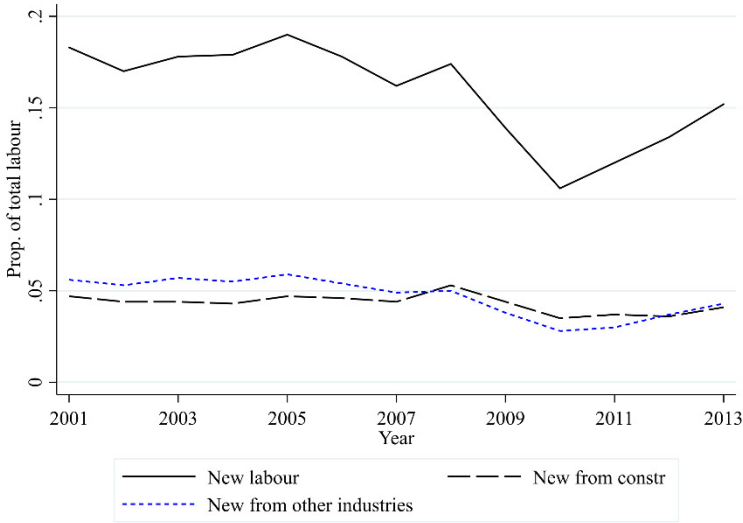
Notes: For construction firms in odd-years from 2005 to 2011 appearing in the BOS innovation module. All regressions also control for year and region fixed effects. The sample is limited to firms sized at least 0.5. Standard errors, in parentheses, are robust and clustered at the firm level. Asterisks denote: *** p<0.01, ** p<0.05, * p<0.10.

Figure 1: Job churn in construction



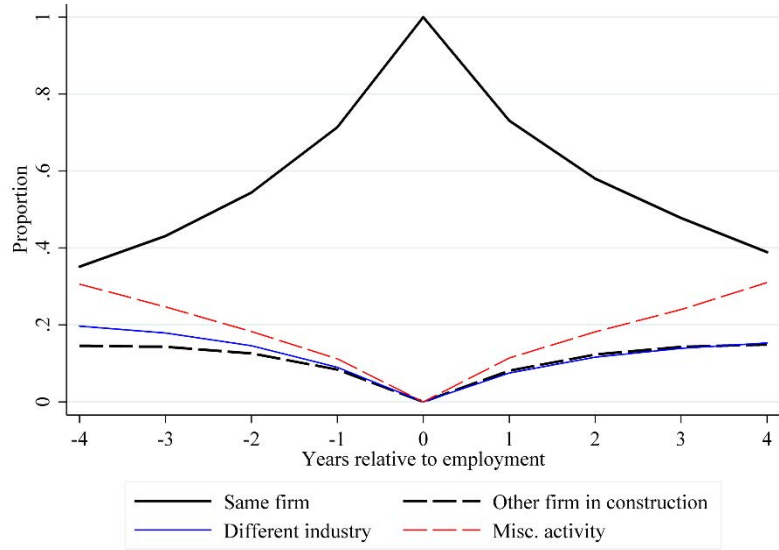
Notes: The job entry rate is the proportion of all construction jobs that are new. The job exit rate is the proportion of all construction jobs that are ending and will not exist again. The job churn rate is the sum of the entry and exit rates.

Figure 2: Proportion of labor in construction that is new




Notes: New labor includes both employee labor (FTE) and working proprietor labor. New labor from a given area (construction or other industries) is calculated as a worker’s current labor input multiplied by the worker’s total labor input last year in that area.

Figure 3: Past and future work of construction workers
 (a) All workers in construction



(b) All new workers in construction entrant firms



Notes: Pooled across cohorts of construction workers from 2005-2010. A person must work at least 0.25 FTE (or contribute any working proprietor labor) for a firm to be counted as working for that firm. When a person works multiple jobs, we consider only the primary job with the highest yearly earnings.

Appendix tables

Appendix Table 1: Simple statistics for regression sample (Table 1)

	mean	10th %-ile	median	90th %-ile	std dev	N
New employees' past prod from constr/firm size	0.001	0	0	0	0.032	156,159
New employees' past prod from constr/firm size (non-zero)	0.004	-0.026	0.000	0.031	0.081	24,888
New employees' past prod from other industries/firm size	0.000	-0.001	0	0	0.040	156,159
New employees' past prod from other industries/firm size (non-zero)	-0.002	-0.036	0.000	0.026	0.085	34,305
New WP past prod from constr/firm size	0.003	0	0	0	0.129	156,159
New WP past prod from constr/firm size (non-zero)	0.102	-0.430	0.049	0.791	0.758	4,476
New WP past prod from other industries/firm size	0.000	0	0	0	0.125	156,159
New WP past prod from other industries/firm size (non-zero)	0.051	-0.795	0.050	1.030	1.294	1,455
Any new employees from constr	0.093	0	0	0	0.291	156,159
Any new employees from other industries	0.115	0	0	1	0.319	156,159
Any new WPs from constr	0.026	0	0	0	0.159	156,159
Any new WPs from other industries	0.008	0	0	0	0.091	156,159
Ever an entrant	1.000	1	1	1	0.000	156,159
1st year	0.142	0	0	1	0.349	156,159
2nd year	0.161	0	0	1	0.368	156,159
3rd year	0.158	0	0	1	0.365	156,159
4th year	0.138	0	0	1	0.345	156,159
Final year	0.065	0	0	0	0.247	156,159
Log gross output	11.78	10.34	11.70	13.47	1.37	156,159
Log capital	9.04	7.58	8.98	10.65	1.27	156,159
Log labor	-0.06	-0.69	-0.34	1.20	0.91	156,159
Log intermediate consumption	11.02	9.17	10.93	13.02	1.57	156,159

Notes: This table presents summary statistics for all firms that are ever seen entering and that appear at least 4 times over the period 2003-2012. The data is limited to the years 2003-2012, and to firm-years with production data and with average employment larger than 0.5.

Appendix Table 2: Simple statistics for regression sample (Table 2)

	mean	10th %-ile	median	90th %-ile	std dev	N
New employees' past prod from constr/firm size	0.000	0	0	0	0.028	281,379
New employees' past prod from constr/firm size (non-zero)	0.003	-0.017	0.000	0.020	0.075	45,717
New employees' past prod from other industries/firm size	0.000	0.000	0	0	0.032	281,379
New employees' past prod from other industries/firm size (non-zero)	-0.002	-0.028	0.000	0.020	0.072	64,500
New WP past prod from constr/firm size	0.001	0	0	0	0.051	281,379
New WP past prod from constr/firm size (non-zero)	0.082	-0.279	0.042	0.539	0.488	4,539
New WP past prod from other industries/firm size	0.000	0	0	0	0.048	281,379
New WP past prod from other industries/firm size (non-zero)	0.002	-0.592	0.035	0.692	0.773	1,572
Any new employees from constr	0.086	0	0	0	0.280	281,379
Any new employees from other industries	0.113	0	0	1	0.316	281,379
Any new WPs from constr	0.010	0	0	0	0.098	281,379
Any new WPs from other industries	0.004	0	0	0	0.059	281,379
Ever an entrant	0.491	0	0	1	0.500	281,379
1st year	0.075	0	0	0	0.263	281,379
2nd year	0.070	0	0	0	0.256	281,379
3rd year	0.069	0	0	0	0.254	281,379
4th year	0.067	0	0	0	0.250	281,379
Final year	0.056	0	0	0	0.230	281,379
Log gross output	11.97	10.40	11.84	13.86	1.48	281,379
Log capital	9.21	7.62	9.10	10.97	1.38	281,379
Log labor	0.11	-0.69	0.00	1.63	1.09	281,379
Log intermediate consumption	11.26	9.34	11.13	13.40	1.64	281,379

Notes: WP stands for working proprietor. This table presents summary statistics for all firms appearing at least 4 times over the period 2003-2012. The data is limited to the years 2003-2012, and to firm-years with production data and with average employment larger than 0.5.

Appendix Table 3: Positive matching between good workers and good firms

High firm fixed effect (>0)	Prop. of workers with high worker FEs (>0)	Observation count
<i>All workers, unweighted</i>		
0	0.376	1,020,639
1	0.413	1,112,346
<i>All workers, fte-weighted</i>		
0	0.411	1,020,639
1	0.449	1,112,346
<i>New workers, unweighted</i>		
0	0.348	402,180
1	0.376	416,079
<i>New workers, fte-weighted</i>		
0	0.385	402,180
1	0.420	416,079

Notes: This table groups firms by whether they are high-paying firms (high firm fixed effect), and then summarizes the skill composition of their work force. It includes employees and working proprietors (WPs), with WPs weighted by pseudo-WP labor in the weighted panels. For all construction firms over the period 2001-2012, with firm size at least 0.5.

Appendix Table 4: Knowledge flows and productivity, extensions and robustness checks

	Log gross output	Log gross output	Log gross output	Log gross output	Log gross output
Extension/robustness check:	<i>Control for new-worker skills, without firm FE</i>	<i>Control for new-worker skills, with firm FE</i>	<i>Weight by gross output, without firm FE</i>	<i>Weight by gross output, with firm FE</i>	<i>New entrants only, without firm FE</i>
New employee mfp at prev constr firms/firm size	0.313***	0.086**	0.224***	0.047	0.318***
	(0.052)	(0.034)	(0.046)	(0.043)	(0.058)
New employee mfp at prev non-constr firms/firm size	0.019	-0.040**	0.094**	-0.023	0.055
	(0.023)	(0.019)	(0.038)	(0.028)	(0.037)
New WP mfp as WP at prev constr firms/firm size	0.128***	0.028**	0.112***	0.038**	0.208***
	(0.014)	(0.012)	(0.012)	(0.015)	(0.025)
New WP mfp as WP at prev non-constr firms/firm size	0.073***	-0.003	0.036*	0.027*	0.112***
	(0.025)	(0.013)	(0.019)	(0.016)	(0.043)
Some new employees from constr	0.001	0.004**	-0.025***	-0.002	0.065***
	(0.003)	(0.002)	(0.008)	(0.005)	(0.010)
Some new employees from other industries	-0.003	0.003*	0.004	0.010**	0.052***
	(0.003)	(0.002)	(0.008)	(0.004)	(0.011)
Some new WPs from construction	0.090***	0.014**	0.017	0.029**	0.122***
	(0.007)	(0.006)	(0.013)	(0.012)	(0.010)
Some new WPs from other industries	0.100***	0.030***	0.029**	0.021*	0.127***
	(0.012)	(0.010)	(0.013)	(0.012)	(0.017)
Avg. skills of new employees/firm size	0.027***	0.009			
	(0.009)	(0.007)			
Avg. skills of new working proprietors/firm size	0.010*	-0.023***			
	(0.006)	(0.006)			
Ever seen entering (2001 onwards)	0.053***		0.024		
	(0.004)		(0.016)		
<i>Observations</i>	281,379	281,379	281,379	281,379	22,203
<i>Number of firms</i>	59,334	59,334	59,334	59,334	21,207
<i>R-squared</i>	0.887	0.664	0.993	0.998	0.799
<i>Firm FE</i>	No	Yes	No	Yes	No
<i>Industries included</i>	All constr with lvl 2 FE	All construction	All construction	All construction	All constr with lvl 2 FE

Notes: Average new worker skills are fte-weighted, and are measured by worker fixed effects and gender-specific age profile parameters in separate wage regressions. All regressions also control for the log of capital, log of labor, log of intermediate materials, year and region fixed effects, and include dummy variables for having production data from the AES survey and being in Canterbury in the post-earthquake period. See further notes in Table 2.

Appendix Table 5: Knowledge flows and productivity, different measure of knowledge flows

	Log gross output	Log gross output	Log gross output	Log gross output	Log gross output	Log gross output	Log gross output	Log gross output
New FTE from productive constr firms/firm size	0.238***	0.047*	0.237***	0.069	0.058	0.159*	0.251***	0.032
	(0.030)	(0.026)	(0.065)	(0.053)	(0.100)	(0.083)	(0.031)	(0.028)
New FTE from unproductive constr firms/firm size	-0.057***	0.003	0.010	0.044	-0.061	0.094	-0.091***	-0.022
	(0.020)	(0.016)	(0.039)	(0.033)	(0.094)	(0.076)	(0.025)	(0.018)
New FTE from productive non-constr firms/firm size	0.040	-0.020	0.083	-0.025	-0.003	-0.057	0.027	-0.017
	(0.032)	(0.027)	(0.063)	(0.041)	(0.114)	(0.120)	(0.038)	(0.034)
New FTE from unproductive non-constr firms/ firm size	0.059**	0.059***	0.105**	0.067*	-0.136	-0.062	0.047*	0.060**
	(0.023)	(0.021)	(0.041)	(0.037)	(0.114)	(0.114)	(0.028)	(0.026)
New FTE from misc sources/ firm size	0.020**	0.012*	0.025*	-0.013	-0.066	-0.019	0.021*	0.026***
	(0.009)	(0.007)	(0.013)	(0.012)	(0.076)	(0.083)	(0.011)	(0.008)
New WP labor from productive constr firms/firm size	0.171***	0.028***	0.146***	0.001	0.267***	0.061	0.179***	0.040***
	(0.010)	(0.009)	(0.017)	(0.016)	(0.053)	(0.051)	(0.012)	(0.011)
New WP labor from unproductive constr firms/ firm size	-0.032***	-0.009	-0.028*	-0.009	0.189*	0.019	-0.039***	-0.012
	(0.010)	(0.011)	(0.016)	(0.020)	(0.109)	(0.070)	(0.013)	(0.012)
New WP labor from productive non-constr firms/ firm size	0.184***	0.041**	0.151***	0.056	0.219**	0.091**	0.194***	0.033
	(0.022)	(0.019)	(0.035)	(0.046)	(0.099)	(0.045)	(0.027)	(0.021)
New WP labor from unproductive non-constr firms/ firm size	-0.044	-0.024	-0.125	0.031	0.207	0.087*	-0.037	-0.052
	(0.041)	(0.029)	(0.109)	(0.047)	(0.149)	(0.045)	(0.038)	(0.038)
New WP labor from misc sources/ firm size	-0.000	-0.039***	0.006	-0.027***	0.087***	-0.008	-0.005	-0.045***
	(0.004)	(0.004)	(0.007)	(0.006)	(0.024)	(0.024)	(0.004)	(0.004)
Ever seen entering (2001 onwards)	0.056***		0.076***		0.127***		0.045***	
	(0.004)		(0.007)		(0.023)		(0.005)	
<i>Observations</i>	281,379	281,379	90,690	90,690	8,577	8,577	182,112	182,112
<i>Number of firms</i>	59,334	59,334	20,025	20,025	1,683	1,683	37,626	37,626
<i>R-squared</i>	0.886	0.670	0.886	0.682	0.942	0.724	0.877	0.660
<i>Firm FE</i>	No	Yes	No	Yes	No	Yes	No	Yes
<i>Industries included</i>	All constr with lvl 2 FE	All constr	Building constr	Building constr	Heavy & civil	Heavy & civil	Constr services	Constr services

Notes: All regressions also control for log of capital, log of labor, log of intermediate materials, year and region fixed effects, and include dummy variables for being an entrant, an exiter, having production data from AES survey, and for being in Canterbury in the post-earthquake period. A productive firm is one whose Cobb-Douglas MFP residual is larger than zero, meaning the firm is more productive than the industry average that year. See further notes in Table 2.

Appendix Table 6: Complete production function regression results, Columns 3 and 4 of Table 2

VARIABLES	Column 3	Column 4
Ever seen entering	0.055*** (0.004)	
New emps' avg mfp at prev constr firms/firm size	0.320*** (0.052)	0.088** (0.035)
New emps' avg mfp at prev non-constr firms/firm size	0.020 (0.024)	-0.039** (0.019)
New WP avg mfp as WP at prev constr firms/firm size	0.128*** (0.014)	0.030*** (0.012)
New WP avg mfp as WP at prev non-constr firms/firm size	0.074*** (0.025)	0.001 (0.014)
Some new employees from constr	0.015*** (0.003)	0.010*** (0.002)
Some new employees from other industries	0.015*** (0.003)	0.010*** (0.002)
Some new WPs from construction	0.090*** (0.007)	0.010* (0.006)
Some new WPs from other industries	0.098*** (0.012)	0.020** (0.010)
Log capital	0.062*** (0.002)	0.094*** (0.003)
Log labor	0.240*** (0.003)	0.195*** (0.003)
Log intermediate materials	0.692*** (0.003)	0.674*** (0.003)
Production data from AES	0.126*** (0.006)	0.088*** (0.005)
Canterbury rebuild dummy	0.095*** (0.008)	0.090*** (0.007)
<i>Observations</i>	281,379	281,379
<i>Number of firms</i>	59,334	59,334
Std. deviation of new emps' mfp at construction	0.064	0.064
Std. deviation of new WP mfp at construction firms	0.756	0.756
<i>R-squared</i>	0.886	0.663
<i>Firm FE</i>	No	Yes
<i>Industries included</i>	All constr with lvl 2 FE	All constr

Notes: See the notes in Table 2. Standard errors, in parentheses, are robust and clustered at the firm level. Asterisks denote: *** p<0.01, ** p<0.05, * p<0.1.

Data Appendix

This study uses rich administrative data from Statistics New Zealand's IDI and LBD. To get comprehensive data on construction firms and their workers, the core of our data comes from the Employer Monthly Schedule (EMS), which shows all employer-employee relationships. The EMS is derived from tax data; each month all employers must file a return with Inland Revenue showing the employees working at the firm that month, the wages paid to them, and the tax deducted.

We use the work of Fabling and Maré (2015a) to identify individuals who are working proprietors at a firm. An individual is counted as a working proprietor if she: pays herself in the EMS (the payer and payee identifiers are the same); reports self-employment income in the IR3 tax form; reports a share of partnership income in the IR7P tax form; or is a company owner receiving payment in the IR4S tax form. Working proprietor relationships are assumed to be permanent, though the labor input from a working proprietor can be zero in a particular year. This identification allows us to separately consider the movements of employees and working proprietors at construction firms. This is important because a large proportion of construction firms in New Zealand are working-proprietor-only firms, and working proprietor relationships may differ on important dimensions compared to employee relationships with firms.

Fabling and Maré (2015a) also provide a measure of employee FTE labor. These measures tend to overestimate true labor input, because they identify FTE as less than 1 only if an employee's income is less than a full-time worker would make on the minimum wage. Despite this, it is better than a simple headcount measure and is the best available labor measure in the IDI.

To these core data we link in numerous other data sources from the IDI, to gain more information on firms and their workers. This is outlined below.

Worker characteristics

Basic demographic information on employees and working proprietors comes from the IDI's core demographic table, which uses various sources for accuracy.²⁰ From this, we capture each worker's age.

We also glean information on whether a worker was recently engaged in formal on-the-job training for trade jobs and apprenticeships. In New Zealand, Industry Training Organisations (ITOs) arrange industry training by working with tertiary education providers. The IDI has comprehensive ITO training data from 2003, showing which individuals were engaged in industry training in a given year and which ITO organized the training.

²⁰ This table is called 'data.personal_detail' in the IDI.

Further information on education comes from the Ministry of Education tertiary education data in the IDI. These data show qualification enrolments and completions (from 1994), as well as course enrolments (from 2001). In documenting the history of construction workers, we simply classify a person as engaging in tertiary study if she studied at least half a full-time equivalent year.

Customs data in the IDI show all border movements into and out of the country, and so allow us to document overseas spells. In documenting the history of construction workers, we classify a person as overseas if he spent at least half the year overseas. Visa data show all visas that were approved, as well as the type of visa. This allows us to identify which workers in construction are recent migrants, and whether they are here on a work visa or a resident visa.

Finally, we use the analysis and data from Maré and Hyslop (2006) and Maré, Hyslop and Fabling (2017) to get measures of observed and unobserved skills for the employees in our data. These come from regressions at the year-job level, where the dependent variable is an employee's log of earnings and the covariates are a set of worker fixed effects; a set of firm fixed effects; and flexible controls for gender-year specific age profiles of earnings. Hence a worker fixed effect represents the portable wage premium wherever a person works, and so constitute a data-driven approach to measuring unobserved skill levels, while the gender-year specific age profiles provide observed measures of skills. These are an important alternative to looking at qualifications, both because qualifications data in the IDI is not comprehensive and because many differences in skill sets between individuals are not captured by differences in qualifications. A person who is a working proprietor can still have a worker fixed effect, if he ever works as an employee for another firm. We do not focus on the skills of working proprietors, but we do consider separately the average skills of new employees and new working proprietors to a firm.

Firm characteristics

Data on firm characteristics come from the LBD, a component of the IDI which contains a variety of tax, administrative, and survey data on all economically active firms in New Zealand.²¹ We use the permanent firm identifiers created by Fabling (2011) to repair broken firm identifiers, which ensures we can track the characteristics and workers of the same firm over time.

Most importantly for this study, each firm in the LBD is assigned a predominant industry, using ANZSIC 2006 codes. These allow us to identify construction firms, as well as worker flows between and within industries. We assign to each firm a permanent industry code, calculated as a firm's industry with the greatest share of employee-months.

²¹ See Fabling and Sanderson (2016) for details on Statistics New Zealand's criteria for an enterprise to be considered economically active.

Sample selection

The broadest estimation sample in this paper is all construction firms in the LBD that have labor input in a given year, for the period 2001-2013. For many descriptive statistics and regressions we limit this to firms in the period 2003-2012, to use firm productivity data and capture the job history of workers new to a firm.

In our productivity analysis, we are limited to firms with useable production data. This decreases the number of firm-year observations from around 550,000 to 360,000 over 2001-2012. As shown by Fabling and Maré (2015b), firms with useable production data in the LBD account for around 60 percent of total industry gross output. In the year ending March 2012, the specific proportions for construction were: 58 percent for 'building construction', 72 percent for 'heavy and civil engineering construction', and 50 percent for 'construction services'. Firms without production data tend to be small, though we cannot say anything about the productivity of firms for which we lack production data.