

From Russia with Labor: The Effects of the Ruble Collapse in an Online Labor Market

*** Preliminary***

Bo Cowgill
Columbia University

John J. Horton
New York University

January 3, 2018

Abstract

The Russian financial crisis and the resulting depreciation of the ruble greatly increased the real returns to Russians from working in online labor markets, as contacts in these market are dollar-denominated and the vast majority of buyers are from the US. We examine the effects of the ruble depreciation in one such online labor market. Russians clearly noticed the opportunity created—the elasticity of total job applications sent by Russians to the price of the USD in rubles was 2.3, with about 80% of the increase explained by Russians entering/staying in the market. Despite a large increase in the real returns to forming a match, we find that Russians did not lower their wage bids by an economically significant amount. On the demand side, employers posting “Russian-compatible” jobs received more applicants as the ruble depreciated, but this influx in applications had no detectable effect on wages, hiring or the posting of additional job openings.

1 Introduction

One of the main ways economists study labor markets is through shocks—ideally exogenous—to labor supply. However, the number of plausibly exogenous shocks is quite limited—immigration or tax policy changes typically provide the best natural experiments, though even in these cases, there are typically interpretive complications.¹ Although researchers can typically measure changes in wages and employment, they have to decide “where to look” in the affected labor market—and market definition is inherently challenging for labor ([Manning and Petrongolo, 2017](#)). The focus in the extant literature on hours and wages is understandable, but relatively little has been said about the effects of supply shocks on the search and matching process, which is central to the “macro labor” perspective ([Mortensen and Pissarides, 1994](#); [Rogerson et al., 2005](#)).

In this paper, we examine the effects of a supply shock in an empirical context that lets us explore “macro labor” considerations, as we observe all applications sent by all workers, the associated wage bids, and the ultimate outcome of that wage bid. The supply shock was a dramatic increase in Russians entering an online labor market during late 2014 and early 2015. What brought them into the market was the sharp decline of the Russian ruble, which fell nearly 60% against the US dollar from January 2014 to March 2015. For Russians, this decline increased the real returns to working in an online labor market, as contracts are dollar-denominated. As there are essentially no employers from Russia in the market, the ruble collapse was not a demand shock. There is no evidence that Russian home-country outside options improved, which would

¹One is the ability to natives to “sort” around new entrants. Another is that firms can increase capital, keeping marginal productivity—and hence wages—the same despite the supply shock. Finally, to the extent new immigrant workers consume, they are also a demand shock. Perhaps unsurprisingly given these complications, the results of this literature are fairly inconclusive—compare [Card \(1990\)](#), which found that the shock of Cuban immigrants from the Mariel boat-lift did not lower wages in Miami, to [Borjas \(2003\)](#), which found that segments of the market that experienced large increases in immigrants did lower wages.

undercut the appeal of working online.

We exploit the fact that the Russian influx was not a general increase in supply, but rather was a shock concentrated in the types of work Russians historically focused on in the market—primarily software development. Country-specific skills specializations are commonplace in our empirical context, and so countries and workers varied in how “exposed” they were to the Russian influx. Critical to our approach is that we can quantify how exposed, *ex ante*, a worker or a collection of workers is to competition from Russians, which we can do because we observe the universe of job applications on the platform i.e., we can see which precisely which other workers a worker was competing with at the level of the job opening.

In our empirical setting, many of the confounding factors that complicate interpretation in conventional settings—general equilibrium effects from new workers also increasing demand, “natives” flowing out or not flowing into the shocked market, adjustments in capital-per-worker—are “off the table” or easily measurable, as we will explain.

In addition to our supply side analysis, we also analyze outcomes for job openings and the reactions of employers. To identify which openings would be affected by the Russian influx, we use historical data to train a machine learning model that take as inputs the title of the job opening, the job description and the required skills. This allows us to create a opening-specific “Russian-compatibility” score, which is the prediction fraction of applicants that would be Russian, in the absence of the ruble collapse. Using our score, we find, as expected, that those openings predicted to receive the most Russian applicants also had the largest influx of Russian applicants after the ruble collapse.

Our main findings focus on the behavior of Russians, non-Russians that compete with Russians, and employers posting Russian-compatible job openings. Following the ruble collapse, Russians, dramatically increased their market participation, as measured by applications sent. As a comparison, we use workers from other countries, controlling for how much workers from those

countries compete with Russians.

In terms of magnitude, we find that a 10% rise in the value of the USD measured in rubles lead to a 23% increase in the number of applications sent per week by Russians. The increase was mostly due to more Russians entering and staying active in the market—which we call an extensive margin effect—rather than active Russians sending more applications. However, the intensive margin effects were non-trivial—a 10% rise in the value of the USD measured in rubles lead to a 5% increase in applications per active Russian worker. We find no evidence that non-Russians exposed to Russian competition exited the market.

If we thought workers' outside options were determined off-platform and offline, then the collapse of the ruble should have greatly lowered the dollar-denominated reservation wage of Russian workers, which in a bargained wage scenario, should lead to lower observed wage bids from Russians. In contrast, if reservation wages are determined on platform—or if workers should be thought of as price-takers—then wage bids should only fall to the extent the Russian influx has an equilibrium effect on the market wage. We can test these predictions in our setting because workers propose hourly wages when they apply to job openings.

We find that Russian workers did not lower their wage bids by an appreciable amount: the point elasticity for the average Russian wage bid with respect to the price of the USD in rubles is just -0.06 ; our 95% confidence interval for the elasticity is $[-0.209, 0.093]$. Non-Russians exposed to Russians competition did not submit detectably lower wage bids. In fact, the point estimate for the most Russia-exposed non-Russian workers is nearly identical to the point estimate elasticity for Russians.

For each application submitted, we observe whether it lead to a hire, and hence we can calculate application win-rates. We might expect that Russians would experience a substantial amount of crowd-out from other Russians, given that those are the workers that Russians most frequently are competing with. However, it is imprecisely estimated and the 95% confidence interval for the

win rate elasticity is $[-1.292, 0.79]$, and hence could clearly contain large effects.

Turning to the demand side, we find that as the ruble collapsed, the most Russian-compatible job openings received more applicants from Russians. For the top 25% of the distribution of Russian-compatible job openings, the elasticity of the number of Russian applications with respect to the price of the USD measured in rubles is 0.73. However, any effect this had on the realized size of the applicant pool was swamped by other source of variation and our estimate of net effects are imprecise one.

We find no evidence that the realized wage bids of those candidates the employer interviewed fell by a detectable amount—though the point estimate is negative. This is consistent with our panel results that showed no evidence that Russians altered their wage bids. We find no evidence that employers were more likely to hire, though we do find strong evidence that they were more likely to hire a Russian. Perhaps not unrelatedly, we find no evidence that employers were more likely to post additional Russian-compatible openings to take advantage of the larger numbers of available Russians. In short, the large increase in Russian applicants had negligible effects even on job openings most likely to be affected by the influx.

The most parsimonious explanation for the pattern of results is that although the increase in Russians was large in relative terms, it was not large enough to move the market. For all job openings posted from the start of our panel in 2014-01-01 until 2014-07-01, Russians made up 1.4% of all hired workers; from 2014-07-01 until 2014-11-01 (during which the ruble began falling sharply), they made up 2.4% of all hired workers. This is a large increase—about 66.7%—but it is still small in absolute terms. Of course, in certain categories of work, the Russian increase was much larger in absolute terms—but even in these sub-markets, we find little evidence for market-level changes in outcomes other than a larger fraction of Russians being hired in response to their increased search efforts.

The rest of the paper is organized as follow. We first discuss the empirical context of our setting in Section 2. Our results are presented in Section 3. Section 4 concludes.

2 Empirical Context

Our empirical context is an online labor market. In these markets, firms hire workers to perform tasks that can be done remotely, such as computer programming, graphic design, data entry, research and writing. Markets differ in their scope and focus, but common services provided by the platforms include maintaining job listings, hosting user profile pages, arbitrating disputes, certifying worker skills and maintaining feedback systems.

Based on dollars spent, the top skills in the marketplace are technical skills, such as web programming, mobile applications development (e.g., iPhone and Android) and web design. Based on hours worked, the top skills are web programming again, but also data entry, search engine optimization and web research, which are non-technical and require little advanced training. The difference in the top skills based on dollars versus hours reflects a fundamental split in the marketplace between technical and non-technical work. There are highly-skilled, highly-paid workers focusing on non-technical jobs, yet a stylized fact of the marketplace is that technical work tends to pay better, generate longer-lasting relationships and require greater skill. These kinds of higher-skill jobs tend to be the focus of Russian workers active on the market.

There has been some research which focuses on the platform marketplace. [Pallais \(2014\)](#) shows via a field experiment that past worker experience in online labor markets is an excellent predictor of being hired for subsequent work. [Stanton and Thomas \(2012\)](#) use online labor market data to show that agencies (which act as quasi-firms) help workers find jobs and break into the marketplace. [Agrawal et al. \(2013\)](#) investigate what factors matter to buyers in making selections from an applicant pool and present some evidence of statistical dis-

crimination; the paper also supports the view of buyers selecting from a more-or-less complete pool of applicants rather than serially screening.

2.1 The collapse of the ruble

The Russian ruble fell nearly 60% against the US dollar from January 2014 to March 2015. This collapse was not due to a general appreciation of the US dollar but rather had Russia-specific causes—namely Russia’s annexation of Crimea and military intervention in Ukraine. The substantial decline in the price of oil during this period is also thought to have contributed.

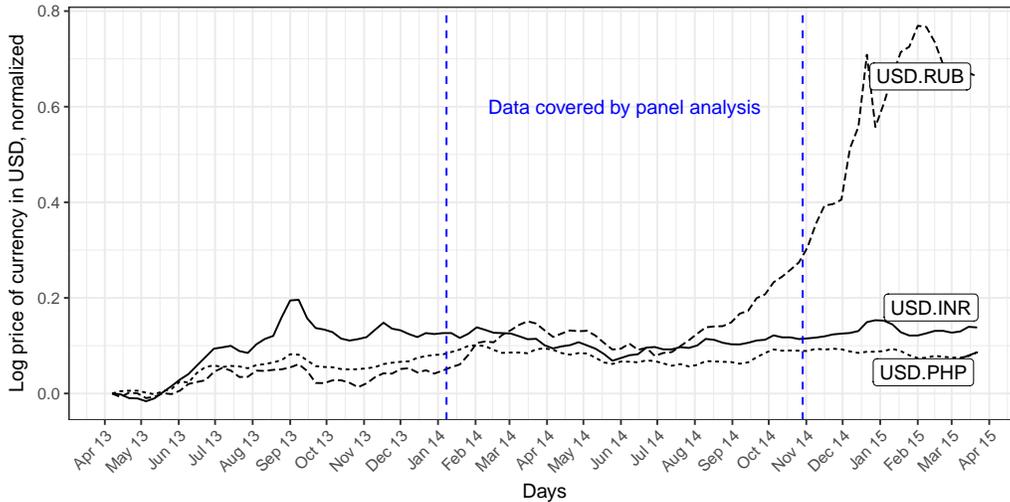
Figure 1 shows the relative prices of Russian ruble two other currencies—the Philippines peso and the Indian rupee—versus the US dollar. All prices are normalized to 1 on April 1st, 2013. We include the rupee and the peso in this plot because those are the “home” currencies of the two largest groups of workers on the platform, namely Filipinos and Indians. For the ruble, we can see that the ruble was more or less constant until about July 2014. By February 2015, the price of the USD was nearly 80% measured in rubles, relative to the start of the figure. The figure also illustrates the period covered by our panel analysis—we will explain these limits in the next section.

2.2 Data

We use several datasets, all created from data obtained from the online labor market in question. The data itself is derived directly from the platform’s own internal database. This database essentially captures nearly everything in the market—jobs posted, applications, wage bids, hires and so on. Critically, we observe the country of the applying worker.

Most of our analysis focuses on the period from 2014-01-08 to 2014-10-29. Although the ruble continued to fall after 2014-10-29, we truncate the sample because at the end of the panel, the platform implemented a market-wide minimum wage, which had large effects on the market equilibrium (Horton, 2017).

Figure 1: Log price (in USD) of the Russian ruble (USD.RUB), Filipino peso (USD.PHP), and Indian rupee (USD.INR) over time



Notes: This figure shows the foreign exchange rate for several currencies with respect to the USD. The rapid depreciation of the ruble can be seen starting in July 2014.

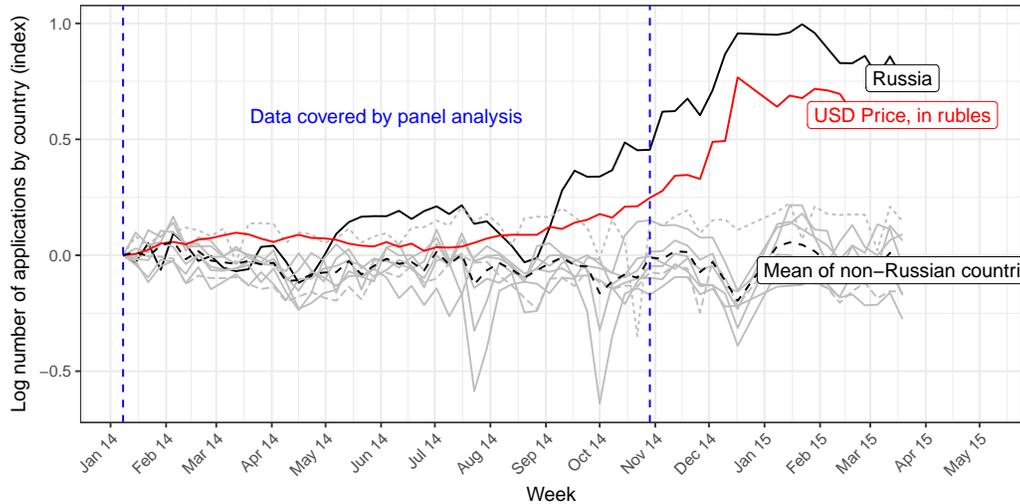
3 Results

To begin, we explore how the collapse of the ruble affected application intensity at the level of worker country, using a country-week panel.

3.1 Job search and matching

The outcomes are measurements that can be defined at the country-week level, such as the number of applications sent, the number of workers active, the number of hires, and so on. For our panel analysis, we use 9 different countries as the comparison units, with Russia included. Using our country-week panel, we plot the log number of applications sent in Figure 2, with each country series demeaned to have a value of 0 at the start of the panel. The series for Russia is shown in a heavy line, whereas all non-Russian countries are plotted in a light line. The average of non-Russian countries is plotted in a dashed

Figure 2: Log number of application sent by week, by country



Notes: This figure shows the time series of the log number of applications by worker, by country, by week for a selection of countries. All time series are demeaned to have a value of 0 in the first week of the panel. The line in red shows the log value of the USD, in rubles over time.

heavy line. In addition to the application intensities, we also plot the value of the USD in rubles over time, which is plotted in red and is also demeaned.

The plot illustrates that before the ruble collapse, Russian workers—like all workers—were responsible for a relatively constant share of applications. However, around April 2014, there is a noticeable increase in applications from Russians, which comes somewhat after a decline in the value of the ruble in February/March 2014. When the ruble begins to increase more strongly September/October, 2014, Russian applications start increasing dramatically as well. At its peak in January 2015, Russians were sending nearly twice as many applications relative to what we would have expected if they followed the average trajectory of other countries.

In Figure 3 we plot the time course for each country separately. We also label Eid Al Fitr (end the end of Ramadan) and the Islamic New Year for 2014 for our two Muslim-majority countries, Bangladesh and Pakistan. We can see a sharp

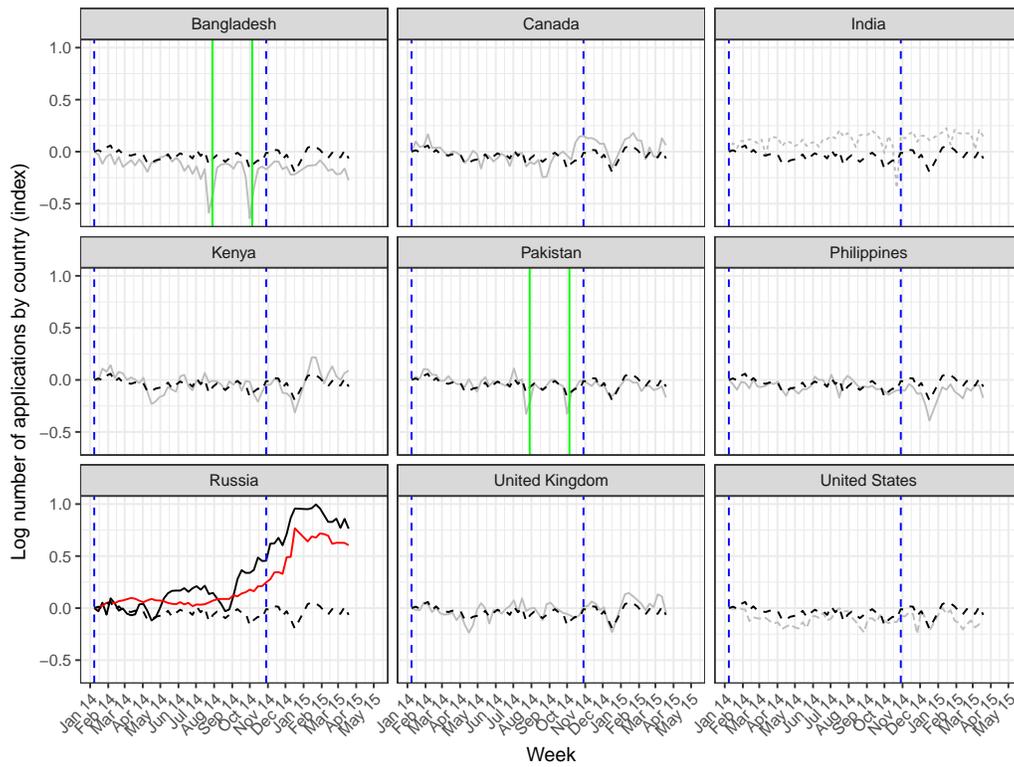
reduction in applications sent around those weeks (which we will control for in our regression analysis).

In addition to total applications, we will also explore the number of workers active in a week as well as their application intensity. We defined “active” defined as sending at least one application. We will also explore the wage bids associated with those applications, their success rate, and finally, the gross number of hires in response to those applications. Although we will explore all of these outcomes in a regression framework, we begin by simply showing graphical evidence. Figure 4 plots the by-country, by-week outcomes for all countries in the panel. The difference series for Russia is indicated in a heavy line. Unlike the previous plots, the various country series are not demeaned.

The panel labeled “Mean wage” illustrates the necessity of truncating the sample in late October 2014, as mean wages jump shortly thereafter for Bangladesh and, not a less extent, Philippines. The cause was the introduction of a platform-wide minimum wage. While the effects are strongly visible for wages, there are noticeable effects for other outcomes of interest, such as win-rates.

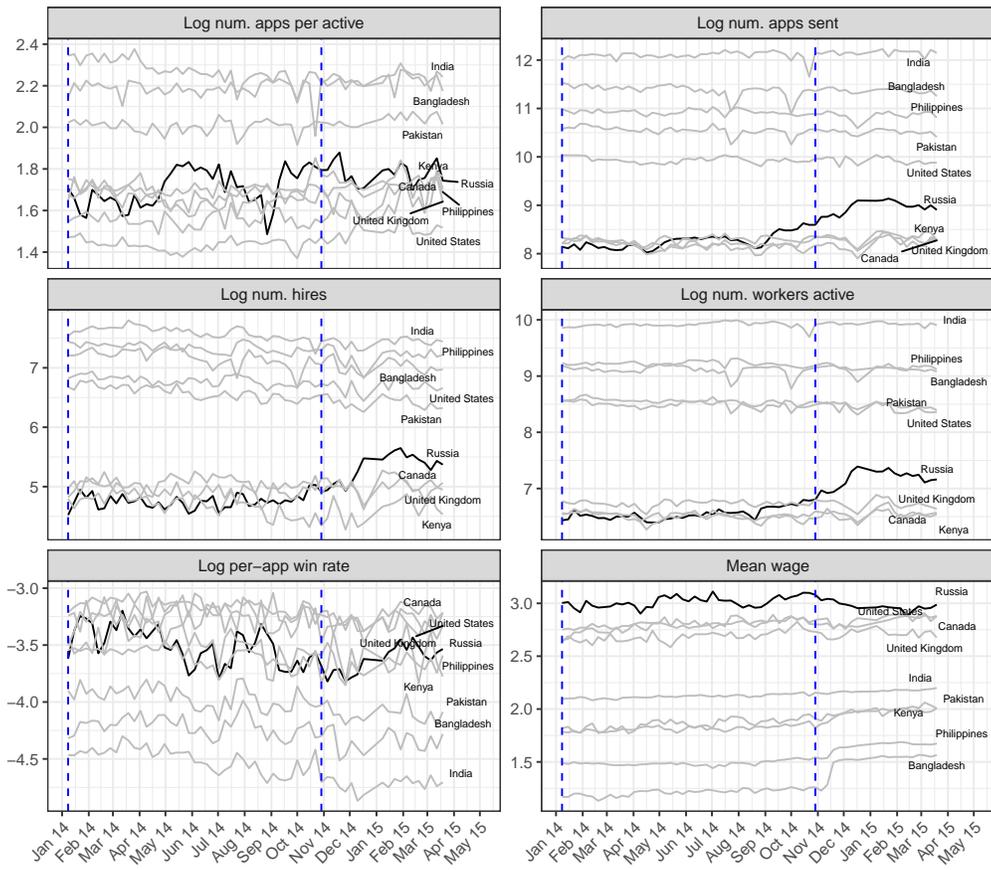
In Figure 5, we plot the same time series as in Figure 4, but with all series demeaned and the non-Russian average plotted. Note that we also put all series on the same scale. As we would expect, this figure previews the main panel regression results. In the top row, we can see in the left panel that application intensity—apps sent conditional upon being active—shows little evidence of expanding as the ruble collapsed. In contrast, the log number of workers active (middle row, right panel) shows a much clearer departure from other countries, as does the total number of applications sent (top row, right panel), as does the number of hires (middle row, left panel). In the bottom row, we see that the Russian time series looks quite similar to the average of other countries, with no strong pattern emerging as the ruble began to depreciate.

Figure 3: Log number of application sent by week by country



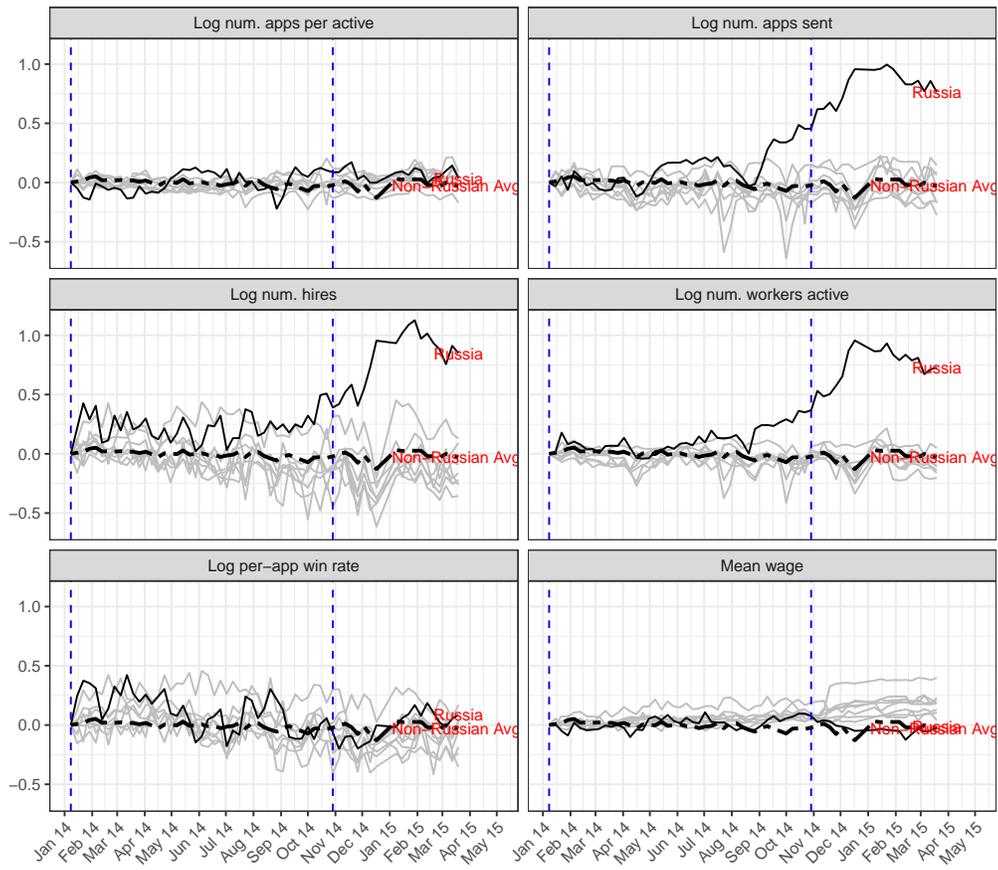
Notes: This figure shows the time series of the log number of applications by worker, by country, by week for a selection of countries. All time series are demeaned to have a value of 0 in the first week of the panel. The line in red shows the log value of the USD, in rubles over time. Eid Al Fitr (end the end of Ramadan) and the Islamic New Year for 2014 for our two Muslim majority countries, Bangladesh and Pakistan, are illustrated with vertical green lines.

Figure 4: Weekly country-level outcomes



Notes: This figure shows the time series of a collection of country-level outcomes.

Figure 5: Weekly country-level outcomes, demeaned



Notes: This figure shows the time series of a collection of country-level outcomes.

3.1.1 Regression evidence

Although the graphical evidence presented in the previous section is suggestive, it does not account for the fact that the ruble collapse may have affected other countries—particularly those workers who compete with Russians. To control for—and explore—this spill-over, we switch to a regression framework.

To begin, we create a country-level measure of how closely workers from that country compete with Russians. This country-specific approach is motivated by a key stylized fact of the marketplace is that workers from the same country tend to specialize in the same kinds of work. As such, a sudden influx of Russians affected certain kinds of work—namely work that Russians tend to specialize in—which in turn could affect certain workers—namely non-Russians that tend to focus on the same kinds of work as Russians.

However, in assessing spill-overs, magnitudes matter and Russians make up a fairly small fraction of all workers in the marketplace. For all job openings posted from the start of the panel in 2014-01-01 until 2014-07-01, Russians made up 1.4% of all hired workers; from 2014-07-01 until 2014-11-01 (during which the ruble began falling sharply), they made up 2.4% of all hired workers. This is a large increase—about 66.7%—but it is still small in absolute terms. Of course, in certain categories of work, the Russian increase was much larger in absolute terms—we will turn to this issue directly later, but for now, we simply focus on country level aggregate measures of exposure.

To create a measure of Russian exposure, we start with each application sent by any worker from any country and compute the fraction of their fellow applicants to that job opening that were Russian. We then average this measure over all workers in a particular country. Using this competition measure, which we turn into a percentile score, s_c^{PCT} , we estimate regressions of the form

$$y_{ct} = \beta_1 (\log p_t \times \text{RUSSIAN}_c) + \beta_2 (\log p_t \times s_c^{PCT}) + \gamma_t + \eta_c + \epsilon \quad (1)$$

where y_{ct} is some outcome of interest, such as log applications sent per week

by workers from country c during week t , p_t is the price of one USD in rubles at the start of week t , RUSSIAN_c is an indicator for whether the observed country c is Russia, and finally, γ_t and η_c are week- and country-specific fixed effects. For Russia, s^{PCT} is set to 0.

Table 1 reports the regression estimates of Equation 1. In Column (1), the outcome is the log number of applications sent—the same outcome plotted in Figure 2. In Column (2), the outcome is the log number of workers active in that country that week. The Column (3) outcome is the log of the number of applications sent divided by the number of workers active—it is essentially an intensive margin estimate of application intensity.²

Starting in Column (1), we can see that a 10% increase in the value of a USD in rubles lead to about 23% more applications from Russians. The interaction term for non-Russian countries is positive, but imprecisely estimated. Column (2) reveals that most of the increase in the applications from Russians was an extensive margin effect: the number of active Russians increased dramatically, with a 10% increase leading to about a 19% increase in active Russians. From the competition exposure interaction term, we can see that there is no evidence that non-Russians highly exposed to competition from Russians were more likely to exit the market—the point estimate is actually positive.

In Column (3), the outcome is the log application intensity. Given the large extensive margin effect relative to the overall increase, there is not much “room” for an intensive margin effect among Russians. Indeed, we see that for each 10% increase in the value of the USD in rubles, Russians only sent about 5% more applications per worker. From the competition index interaction, we have some slight evidence that non-Russians increased their application intensity, but the effects are imprecisely estimated.

²This measure is mechanically the difference between the total active and total applications measures because of the log transformation, but we include it anyway for ease of comparison.

Table 1: Effects of the ruble collapse on worker job search, by country

	<i>Dependent variable:</i>		
	Log # apps sent (1)	Log # workers active (2)	Log apps/active (3)
$\log p_t \times \text{RUSSIAN}_c$	2.323*** (0.136)	1.863*** (0.134)	0.460*** (0.129)
$\log p_t \times s_c^{PCT}$	0.513 (0.369)	0.299* (0.181)	0.213 (0.281)
Observations	387	387	387
R ²	0.997	0.998	0.974
Adjusted R ²	0.997	0.998	0.970
Residual Std. Error (df = 333)	0.081	0.056	0.048

Notes: This table reports the results of regression estimates of Equation 1. The data consists of country-week observations of application behavior by workers. In Column (1), the outcome is the log number of applications sent by workers from a particular country that week. The main independent variable is p_t , which is the price of 1 USD, measured in Russian rubles at the start of week t . This is interacted with RUSSIA_c , which is an indicator if the country in question is Russia. It is also interacted with s_c^{PCT} , which is a country's percentile in a measure of how exposed workers from that are to competition from Russian workers (which is set to 0 for Russia). The outcome in Column (2) is the log number of workers in a country active that week, defined as sending at least one application. The outcome in Column (3) is the log application intensity, with intensity being the number of applications sent per active worker. Standard errors are clustered by country. Significance indicators: $p \leq 0.10$: †, $p \leq 0.05$: *, $p \leq 0.01$: ** and $p \leq .001$: ***.

3.1.2 Wage bidding, per-application win rates, and total hires

We now turn to measures of worker bidding behavior and application outcomes. Table 2 reports regressions with the same specification as Equation 1. In Column (1), the outcome is the log average wage bid. In Column (2), the outcome is the log application “win rate” or the log of the number of hires divided by the number of applications sent. In Column (3), the outcome is the log number of hires.

Starting with wage bids, although the effect for Russians is negative, the effect is small in magnitude. A 10% in the USD measured in rubles lowered Russian wage bids on average by about -0.58% . However, this elasticity is imprecisely estimated—the 95% confidence interval for the elasticity is $[-0.209, 0.093]$. Non-Russian countries exposed to Russian competition showed no change in their wage bids—the point estimate is close to zero, though the effect is also imprecise.

For application win-rates, Column (2) shows that the effect of the ruble collapse is imprecisely estimates. The 95% confidence interval for the elasticity is $[-1.292, 0.79]$. For non-Russians, there is no evidence of crowd-out, as the point estimate is positive.

In Column (3), the outcome is the log number of hires. We can see that Russians were hired more frequently and that the effect is sizable. The hire elasticity with respect to the USD price in rubles is 2.1. This is smaller than the application elasticity of 2.3, which reflects the lower per-application win-rate for Russians.

3.2 Country-specific skill specialization

The increase in the number of applications from Russians was not evenly spread across all kinds of work. Figure 6, shows how the fraction of applicants from Russia changed over time. The right panel, labeled “collapse year”, shows the fraction of applications during two periods: The pre-period is April 1, 2014 to

Table 2: Effects of the ruble collapse on wage bidding and win rates

	<i>Dependent variable:</i>		
	Log mean wage bid	Log app win rate	Log # hires
	(1)	(2)	(3)
$\log p_t \times \text{RUSSIAN}_c$	-0.058 (0.075)	-0.251 (0.521)	2.072*** (0.611)
$\log p_t \times s_c$	-0.046 (0.152)	0.042 (0.656)	0.555 (0.880)
Observations	387	387	387
R ²	0.998	0.966	0.994
Adjusted R ²	0.998	0.960	0.993
Residual Std. Error (df = 333)	0.029	0.091	0.097

Notes: This table reports the results of regression estimates of Equation 1. The data consists of country-week observations of application behavior and outcomes. The main independent variable is p_t , which is the price of 1 USD, measured in Russian rubles at the start of week t . This is interacted with RUSSIA_c , which is an indicator if the country in question is Russia. It is also interacted with s_c , which is a measure of how exposed a worker from country c is to competition from Russian workers. The outcome in Column (1) is the log win rate, defined as the log of total hires divided by total applications sent. In Column (2), the outcome is the average log wage bid from workers from that particular country that week. Standard errors are clustered by country. The outcome in Column (3) is the log number of “hires” defined as applications that lead to an accepted job offer to the worker submitting that application. Significance indicators: $p \leq 0.10$: †, $p \leq 0.05$: *, $p \leq 0.01$: ** and $p \leq .001$: ***.

July 1st, 2014, which is indicated by a circle, and a post period from January 1st, 2015 to March 1st, 2015, which is indicated by triangle. The left panel, labeled “1 year before collapse year” is the same as the right panel, but with the observations shifted back one year. In collapse year, in both the pre- and post-periods (in the right panel) and the same comparisons pushed one calendar year in the past (in the left panel). A circle corresponds to the “pre” period and the triangle to the “post” period.

We can see that in certain categories and sub-categories, the collapse of the ruble lead to large increases in the fraction of applications coming from Russians, whereas in other areas, there was minimal impact. The plot makes it clear that the influx of Russians was concentrated in categories of work that Russians already focused on before the ruble collapse.

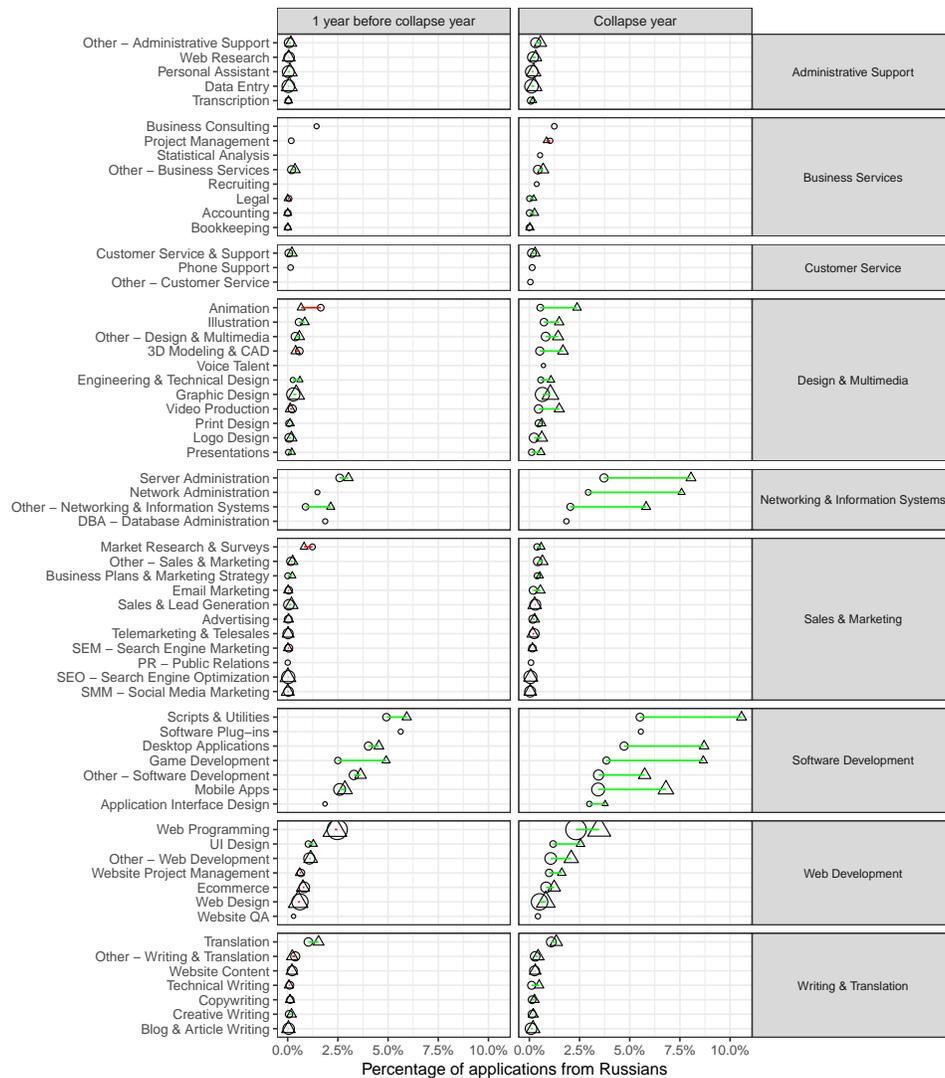
3.3 Outcomes for Russian-compatible vacancies

Employers posting jobs that would appeal to Russians presumably received more applications and lower wage bids—this is the main thrust of the results from the individual worker analysis. To see how this affected the employer’s decision, we need to switch to the job post, or vacancy, as our unit of analysis. A key complication is that we know need some measure of how appealing a job opening would be to Russian applicants.

To construct a measure of Russian exposure, we first construct a historical dataset of job openings, recording the fraction of those applicants that were Russian. Then, we use the full document-term-matrix for the skills required for that opening, as well as another of other characteristics set by the employer. We then use gradient boosting (using the `xgboost` R package) to train a model a linear model (Friedman et al., 2000; Chen et al., 2015). We use the fitted model out of sample to make predictions for all vacancies posted during our panel year. The predictive model can currently explain about 25% of the variance in the realized fraction of Russian applicants in our panel sample.

Figure 7, in the right panel, we plot the mean predicted score (the fraction

Figure 6: Changes in the fraction of applicants from Russia



Notes: This figure shows fraction of applicants from Russia, per the platform sub-category. The right panel, labeled “collapse year”, shows the fraction of applications during two periods: The pre-period is April 1, 2014 to July 1st, 2014, which is indicated by a circle, and a post period from January 1st, 2015 to March 1st, 2015, which is indicated by triangle. The left panel, labeled “1 year before collapse year” is the same as the right panel, but with the observations shifted back one year. Point sizes are scaled by the number of total job openings in that sub-category. The sample is restricted to categories with at least 300 total openings in both pre and post periods for the collapse year and the year before the collapse.

of Russian applicants), by quantile (with the cut points determined by pooling over the entire panel). We can see that the predicted score shows no time trend in any level—consistent with there being no demand shift that could explain the pattern.

In the left panel, we plot the mean actual, realized fraction of Russian applicants, but by predicted quantile. At first, the predicted and actual match during the period before the ruble began to strongly depreciate. However, as the ruble depreciates, we can see somewhat of an increase in every quantile—even the lowest—there is an increase in the fraction of actual applicants, though the increase is concentrated among those job openings expected to receive the largest number of Russian applicants. Among the top 25%, the fraction nearly doubles.

To see how this Russian supply shock affected the openings, we discretize the Russian compatibility score in a collection of k percentile intervals, from lowest score to highest score. We then estimate regression of the form

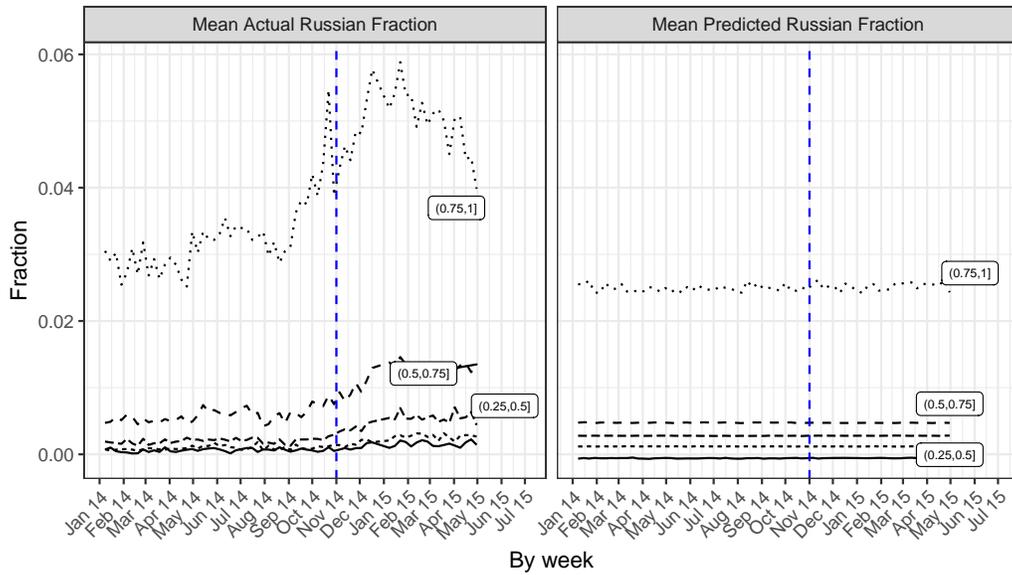
$$y_j = \sum_k \alpha_k \text{RCSINT}_{k(j)} + \sum_k \beta_k (\text{RCSINT}_{k(j)} \times \log p_{t(j)}) + \gamma_{t(j)} + \epsilon \quad (2)$$

where y_j is some outcome for job opening j , $\text{RCSINT}_{k(j)}$ is an indicator for whether the Russian-compatibility score for job opening j is in the k th interval, and $p_{t(j)}$ is the price of 1 USD in rubles when opening at time t , which is when opening j was posted, and finally, γ_t is day fixed effect.

Table 3 reports the $\hat{\beta}_k$ regressions where the outcomes are various measures of the number of applicants received. In Column (1), the outcome is the log number of Russian applications received, plus 1. For vacancies with relatively lower Russian-compatibility scores, the effect is positive but not large in magnitude. However, the effect is increasing in the Russian compatibility score. At the highest band—the most Russian compatible openings—the increase is large and the elasticity is 0.73.

How this Russian influx translated into the size of the actual applicant pool

Figure 7: Comparison of the predicted fraction of Russian applicants (based on job opening characteristics) versus the realized fraction



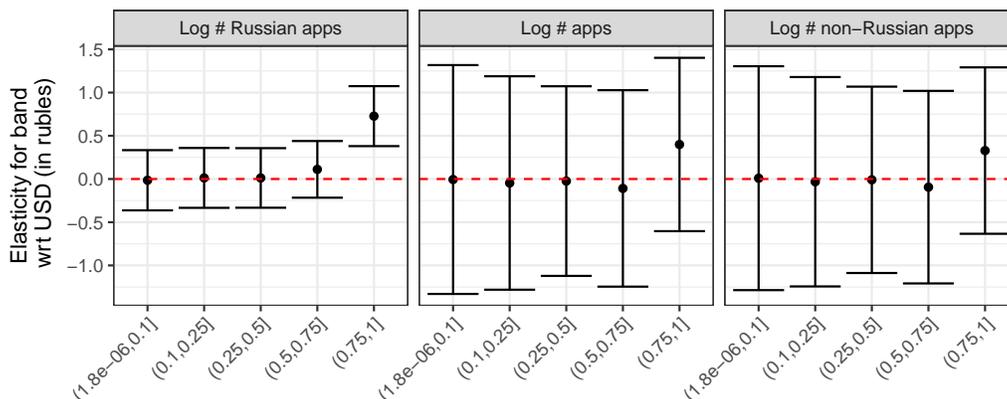
Notes: This figure shows the predicted and actual fraction of Russian applicants for job openings, with the predictions derived from the text of the job description, skill required and job title. The different lines correspond to the different quantiles of the predicted Russian applicant fraction distribution.

Table 3: Effects of the ruble collapse on per-vacancy measures of competition, wage bidding, hires and total wage bill

	<i>Dependent variable:</i>		
	Log # Russian apps	Log # apps	Log # non-Russians apps
	(1)	(2)	(3)
$\log p_{t(j)} \times \text{RCSINT} = (1.8e - 06, 0.1]$	-0.014 (0.174)	-0.006 (0.662)	0.010 (0.648)
$\log p_{t(j)} \times \text{RCSINT} = (0.1, 0.25]$	0.013 (0.173)	-0.046 (0.617)	-0.032 (0.606)
$\log p_{t(j)} \times \text{RCSINT} = (0.25, 0.5]$	0.012 (0.172)	-0.024 (0.549)	-0.009 (0.539)
$\log p_{t(j)} \times \text{RCSINT} = (0.5, 0.75]$	0.112 (0.164)	-0.109 (0.568)	-0.094 (0.557)
$\log p_{t(j)} \times \text{RCSINT} = (0.75, 1]$	0.727*** (0.173)	0.399 (0.501)	0.329 (0.481)
Day FE	Y	Y	Y
Mean outcome (in levels)	0.3	31.81	31.51
Observations	303,104	303,104	303,104
R ²	0.217	0.007	0.008
Adjusted R ²	0.217	0.007	0.008

Notes: This table reports regressions on vacancy level outcomes on changes in the price of the USD (measured in rubles) and its interaction with indicators how Russian-compatible that opening is, based on a model fit with historical data. Significance indicators: $p \leq 0.10$: †, $p \leq 0.05$: *, $p \leq 0.01$: ** and $p \leq .001$: ***.

Figure 8: Elasticities of per-application measures of supply with respect to the price of the USD, in rubles

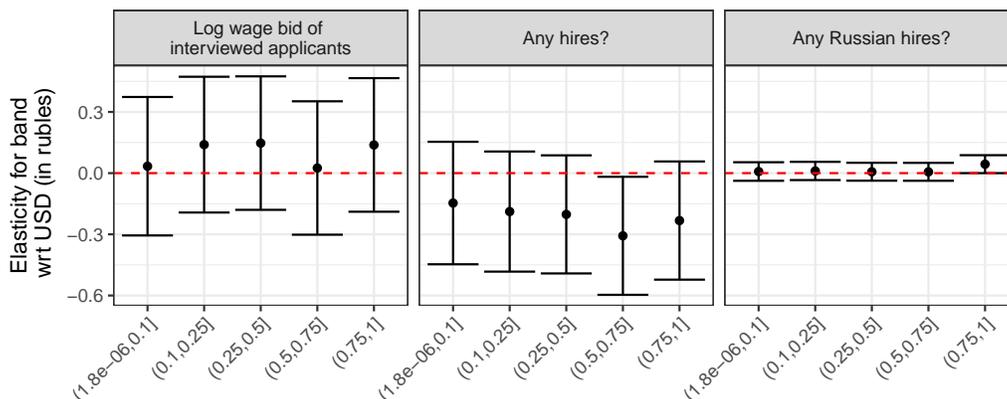


Notes: This figure plots the β_k coefficients from an estimate of Equation 2 for several outcomes.

is considered in Column (2), which reports the results of a regression where the outcome is the log number of applicants. Although the point estimates are substantial for some intervals, they are highly imprecise. Similarly, in Column (3), the outcome is the log number of non-Russian applicants, plus 1 and again, the point estimates are imprecise. In short, there is little that can be said about the effects of Russian entrants on applicant pool sizes. As it will be more informative to view the collection of elasticities for other outcomes graphically, in Figure 8, we re-report the coefficients from Table 3.

Now we turn to wage bidding and hiring. In Figure 9, we turn to the effects of the ruble collapse on measures of wage bids and hiring at the job opening level. In the leftmost panel, the outcome is the average log wage of interviewed candidates. There is no evidence of that the ruble collapse had any effect on the wage bid of interviewed applicants, even among the most Russian-compatible openings. In the next panel, we look at whether the employer hired anyone at all. The point estimates are all centered around zero and there is no evidence that more Russian-compatible openings had higher hiring rates.

Figure 9: Elasticities of per-application measures of wage bids and hiring with respect to the price of the USD, in rubles



Notes: This figure plots the β_k coefficients from an estimate of Equation 2 for several outcomes.

In the rightmost panel, the outcome is whether the employer hired a Russian. Here we see positive effects from about the 75th percentile onward, and large effects at the highest percentile. In short, at the level of the applicant, even for job openings most exposed to the Russian influx, the effects on realized wage bids and hiring was non-existent. The only finding that seems fairly clear is that more Russians were ultimately hired among the most Russian compatible job openings.

4 Conclusion

We report the results of shock that increased the real returns to working for some workers. In response to the increased value of a match on the platform, Russians sent more applications. This was primarily an extensive margin response, in that previously inactive Russians became active, though active Russians did send somewhat more applications. We have no evidence that they altered their wage bids at all. In terms of competition, we find that non-Russians

did not leave the market.

Employers posting the kinds of jobs that Russians specialized in enjoyed somewhat more applications after the ruble collapse. However, they were apparently no more likely to hire. Furthermore, we find no evidence that employers were more likely to post additional openings to take advantage of the larger applicant pools.

References

- Agrawal, Ajay K., Nicola Lacetera, and Elizabeth Lyons**, “Does information help or hinder job applicants from less developed countries in online markets?,” January 2013, (NBER Working Paper 18720).
- Borjas, George J**, “The Labor Demand Curve Is Downward Sloping: Reexamining the Impact of Immigration on the Labor Market,” *The Quarterly Journal of Economics*, 2003, 118 (4), 1335–1374.
- Card, David**, “The impact of the Mariel boatlift on the Miami labor market,” *Industrial & Labor Relations Review*, 1990, 43 (2), 245–257.
- Chen, Tianqi, Tong He, Michael Benesty et al.**, “Xgboost: extreme gradient boosting,” *R package version 0.4-2*, 2015, pp. 1–4.
- Friedman, Jerome, Trevor Hastie, Robert Tibshirani et al.**, “Additive logistic regression: a statistical view of boosting (with discussion and a rejoinder by the authors),” *The annals of statistics*, 2000, 28 (2), 337–407.
- Horton, John J**, “Price floors and employer preferences: Evidence from a minimum wage experiment,” 2017.
- Manning, Alan and Barbara Petrongolo**, “How local are labor markets? Evidence from a spatial job search model,” *American Economic Review*, 2017, 107 (10), 2877–2907.
- Mortensen, Dale T and Christopher A Pissarides**, “Job creation and job destruction in the theory of unemployment,” *The Review of Economic Studies*, 1994, 61 (3), 397–415.
- Pallais, Amanda**, “Inefficient hiring in entry-level labor markets,” *American Economic Review*, March 2014, (18917).

Rogerson, Richard, Robert Shimer, and Randall Wright, “Search-theoretic models of the labor market: A survey,” *Journal of economic literature*, 2005, 43 (4), 959–988.

Stanton, Christopher and Catherine Thomas, “Landing the first job: The value of intermediaries in online hiring,” *Available at SSRN 1862109*, 2012.

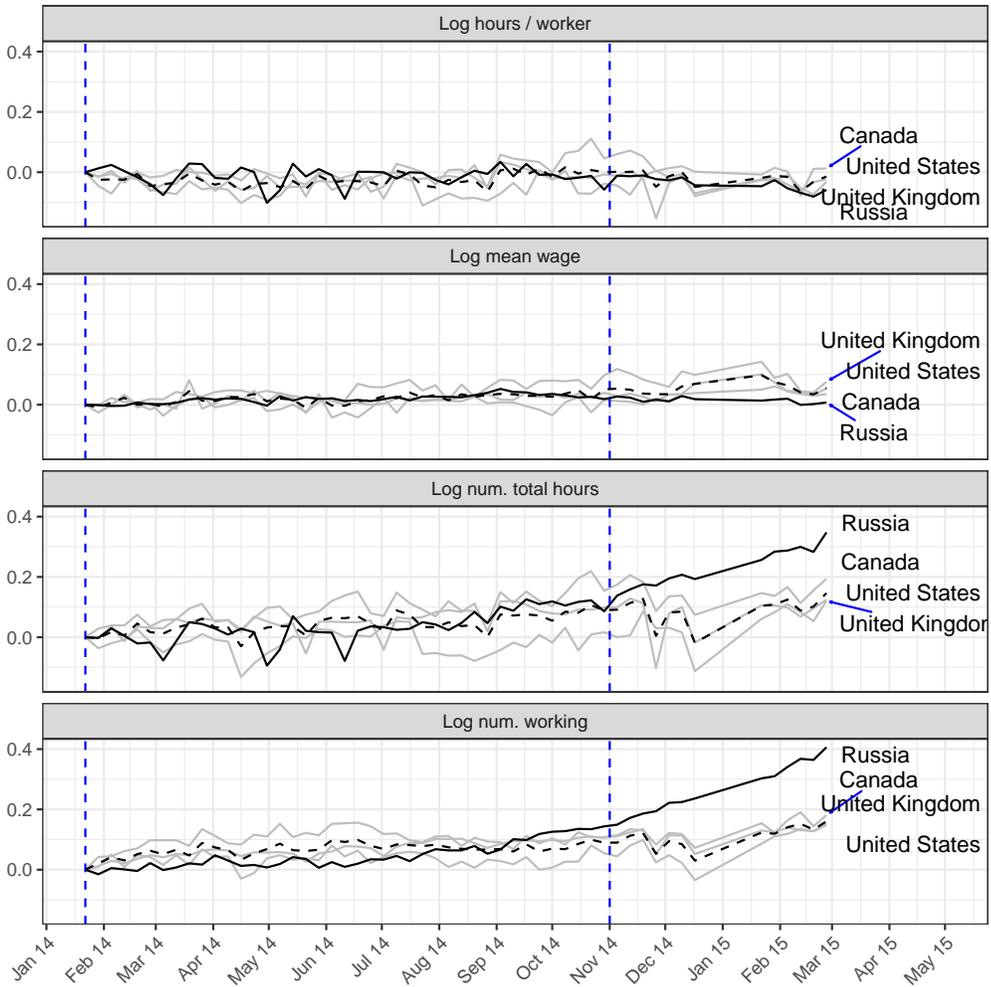
A Hours-worked

After a contract is formed for hourly work, the platform intermediates the relationship, recording hours-worked using a proprietary time-tracker that workers install on their computers. We use these time reports to construct a week-country panel of hours-worked, number of workers working some number of hours, and at what average hourly week. In Figure 10, we plot the log value of the various panel outcome measures by country, over time. We select countries that had broadly similar characteristics to Russia (focusing on relatively high-end, high-wage work) to allow for more credible comparisons in the post-minimum wage period.

The plot shows that the number of workers working and the total hours-working clearly increased, though most the increase appears concentrated in the period after the minimum wage imposition (which coincidentally is also when the ruble fell the most). The relative lack of effects in September/October 2014—despite large increases in applications and hires during this period—is consistent with job-finding (and hence hours-worked) taking time to form.

Interestingly, the plot shows no evidence of an intensive margin effect, in the sense that hours-worked per worker is similar to other comparison countries. There is some evidence of a decline in wages, though the effect appears to be quite small. This plot alone cannot distinguish between composition effects and within-worker changes in wage bidding. For this composition versus behavior question, we will switch to a regression framework with individual data.

Figure 10: Country-week panel showing the number of workers working some number of hours, total hours-worked, hours-per-worker and average wages



Notes: This figure shows country-week panel means for measures of labor supply and hourly wages.

Before breaking into individual level panels, we re-capitulate the analysis in Figure 10 in regression form, in Table 4. In Column (1), the outcome is the log number of workers active, here defined as working some number of hours. This can be thought of as the extensive margin labor supply elasticity (to the platform). In Column (2), the outcome is the log number of hours-worked in total. Finally, in Column (3), the outcome is the log number of hours worked per active worker.

Table 4: Effects of the ruble collapse on labor supply

	<i>Dependent variable:</i>		
	Log # active working (1)	Log hours-worked (2)	Log hours/active (3)
$\log p_t \times \text{RUSSIAN}_c$	0.486*** (0.008)	0.397*** (0.021)	-0.089*** (0.015)
Country FE	Y	Y	Y
Week FE	Y	Y	Y
Observations	252	252	252
R ²	0.914	0.976	0.977
Adjusted R ²	0.883	0.968	0.969

Notes: Significance indicators: $p \leq 0.10$: †, $p \leq 0.05$: *, $p \leq 0.01$: ** and $p \leq .001$: ***.

For Column (1), we find an extensive margin elasticity of 0.5. This is smaller than the total hours elasticity in Column (2), which is just 0.4. Unsurprisingly given this difference, in Column (3), we see a negative effect on hours per worker. However, as we will see in the next section, these aggregate effects reflect composition effects. In a nutshell, there was a large extensive margin effect which brought in “new” Russians who work relatively few hours, perhaps involuntarily, as new entrants have lower job-finding probabilities. When we restrict ourselves to a within-worker analysis, we find, as expected, a positive intensive margin elasticity.

B Labor supply panel evidence at the individual level

We create an unbalanced worker-week panel of Russians and the same non-Russians used in Section A. With this panel, we only observe a worker-week if they worked some amount of hours. We report regressions of log hours-worked and log average wages per week, with and without worker-specific fixed effects. These regressions are reported in Table 5. In Columns (1) and (2), the outcome is the log average hourly wage. In Column (1), we do not include a worker-specific fixed effect, whereas in Column (2) we do. Both specifications have weekly time fixed effects.

Table 5: Effects of the ruble collapse on wages and hours

	<i>Dependent variable:</i>			
	Log avg. wage		Log hours	
	(1)	(2)	(3)	(4)
$\log p_t \times \text{RUSSIAN}_c$	-0.175*** (0.024)	-0.030*** (0.011)	-0.064 (0.041)	0.188*** (0.037)
Worker FE	N	Y	N	Y
Russian FE	Y	N	Y	N
Observations	365,979	365,979	365,979	365,979
R ²	0.020	0.956	0.201	0.684
Adjusted R ²	0.019	0.953	0.200	0.662

Notes: Significance indicators: $p \leq 0.10$: †, $p \leq 0.05$: *, $p \leq 0.01$: ** and $p \leq .001$: ***.

The important of worker-specific effects is clear—without them, we get average wages declining in the price of the USD in rubles by about 1.8% for each 10% increase, whereas with worker fixed effect, the magnitude is only a quarter as large. This between versus within distinction could likely reflect that is an unbalanced panel, Russians joining later may have lower reservation wages and/or be less experienced, which shows up in the cross-section, but not the panel.

This hypothesis that composition effect matter is also consistent with the pattern of results for hours-worked, in Columns (3) and (4). In cross-sections, we see a point estimate close to zero for the interaction term and of the “wrong” sign. Switching to the within-worker evidence, we find a positive effect, with an implied intensive margin elasticity of 0.2. If “new” Russians have trouble breaking into the market and obtaining jobs, then their hours-worked is likely to be smaller, at least at first.

Now we switch to a balanced panel where our only outcome of interest is an indicator for whether the worker worked some number of hours that week. Table 6 shows that regardless of the including of worker fixed effects, as the ruble began to depreciate, Russians were more likely to work at least some hours. The effects are substantial—the baseline probability of working some week with this full panel is just about 25%, so a 10% increase in the USD having a 0.017 effect in levels is almost 10%. However, this is somewhat misleading as some late Russians mechanically could not have been active early in the panel.

Table 6: Effects of the ruble collapse whether the worker worked any hours in a week

	<i>Dependent variable:</i>	
	Any hour worked?	
	(1)	(2)
$\log p_t \times \text{RUSSIAN}_c$	0.177*** (0.013)	0.177*** (0.013)
Worker FE	N	Y
Russian FE	Y	N
Observations	1,464,120	1,464,120
R ²	0.010	0.441
Adjusted R ²	0.010	0.432
Residual Std. Error	0.431 (df = 1464055)	0.326 (df = 1440817)

Notes: Significance indicators: $p \leq 0.10$: †, $p \leq 0.05$: *, $p \leq 0.01$: ** and $p \leq .001$: ***.

We now turn to considering only those workers that were active “early” in the panel, before the ruble fell. In Table 7, the sample is restricted to those workers who worked at least once before 2014-03-01.

In Column (1), the outcome is an indicator for whether the worker worked that week. An increase in the value of the USD does make it more likely that a Russian worker works that week. However though it is smaller than the effect from the full panel, which is unsurprising in that this “early” panel can only capture workers choosing not to exit.

In Column (2), the outcome is the log average wage. The coefficient on the interaction term is a precisely estimated zero—there is no evidence that already-working Russians lowered their wage bids as the ruble began to fall. However, Column (3) makes it clear that they began working more hours, conditional upon working any—for a 10% increase in the value of the USD, they worked about 2% more hours per week.

Table 7: Effects of the ruble collapse on participation, wages and hours-worked for pre-ruble collapse incumbent workers

	<i>Dependent variable:</i>		
	Any hours?	Log avg. wage	Log hours-worked
	(1)	(2)	(3)
$\log p_t \times \text{RUSSIAN}_i$	0.092*** (0.019)	-0.017 (0.012)	0.189*** (0.043)
Worker FE	Y	Y	Y
Conditional on working	N	Y	Y
Observations	498,645	227,726	227,726
R ²	0.540	0.956	0.683
Adjusted R ²	0.533	0.954	0.671

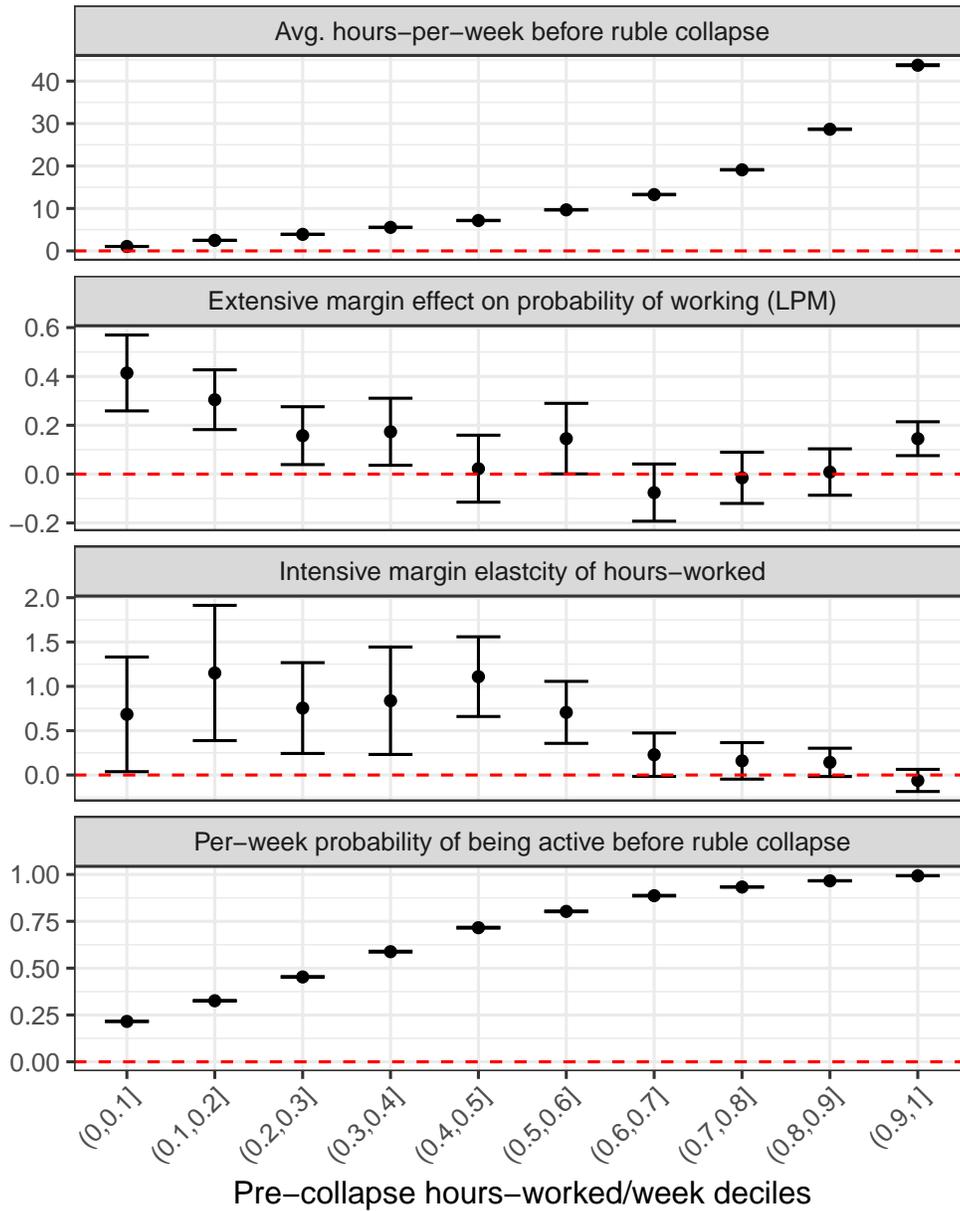
Notes: Significance indicators: $p \leq 0.10$: †, $p \leq 0.05$: *, $p \leq 0.01$: ** and $p \leq .001$: ***.

B.1 Intensive margin labor supply elasticity by current hours-per-week

Incumbent workers would likely respond differently on the intensive margin depending on how many hours they are already working. Russian workers that were already working 40+ hours per week would have little capacity to take on more work. To assess whether intensive margin elasticities vary by pre-collapse hours-worked, we take all incumbent workers and compute their pre-collapse average hours of work per week. Figure 11, in the top panel reports the averages by decline. We can see that the median active worker was only working about 10 hours per week. In contrast, workers in the top decline were working more than 40 hours per week.

To compute elasticities, we interact these decile indicators with the log USD price and the Russian individual indicator. We do condition on

Figure 11: Hours-worked and estimated intensive margin elasticities for workers active before the ruble collapse



Notes: