Subsidizing Job Search Considered Harmful: Evidence from a Field Experiment (Preliminary)*

Apostolos Filippas
Fordham

Andrey Fradkin
BU

John J. Horton
MIT & NBER

July 25, 2023

Abstract

New entrants to a labor market had their job search efforts subsidized, with the subsidy amount varying randomly. The subsidy was highly effective at increasing job search intensity and subsequent average earnings. At the best-performing subsidy level, \$1 in subsidy leads to \$16 in additional job-seeker earnings, even accounting for crowd-out. However, the increase in the non-financial cost of job search more than dissipates the welfare benefits, making the subsidy inefficient unless the social planner puts a very high weight on novice worker earnings—which they might want to do.

JEL Codes: A11, B22, C33

*Author contact information, code, and data are currently or will be available at `http://www.apostolos-filippas.com/`.

1 Introduction

Should job search efforts be subsidized? Or, in the language of labor market matching models, would the additional search effort from subsidization contribute enough to the formation of matches to be worth the cost (Mangin and Julien, 2021; Hosios, 1990)? There are some practical and economic challenges to efficient subsidization. A policy that successfully increased search effort might simply crowd-out other job-seekers (Crépon, Duflo, Gurgand, Rathelot and Zamora, 2013; Barach, Golden and Horton, 2020)—though net increases in match formation are possible (Horton, 2017a). Crowd-out aside, if job-seekers are already making the optimal choice about search intensity given constraints, loosening their budget constraint by a transfer will increase their welfare in expectation only by the transfer amount. However, there are also good reasons to think job-seekers might make sub-optimal choices. They could be badly informed about the returns to additional search (Miano, 2023; Beam, 2021; Jäger, Roth, Roussille and Schoefer, 2021), face credit constraints, behavioral biases (DellaVigna, Heining, Schmieder and Trenkle, 2021), be inefficiently risk-averse, and so on. This characterization might be particularly true for novice job-seekers less familiar with the market.

In this paper, we report the results of a field experiment that subsidized job search efforts for new entrants to an online labor market. Job seekers applying to jobs hosted on the platform must use tokens, which we call "coins." While some coins are given when workers first join the platform, they must buy additional tokens from the platform once they exhaust their initial grant, creating a financial cost to workers from applying to a job. The subsidization experiment took the form of different-sized grants to workers when they joined the platform. The platform created a marginal financial cost of a job application to address congestion issues that arise when the technical cost of a job application is low.

The subsidy has several unique and useful economic properties. The subsidy was purely financial. As such, it is possible to directly and unambiguously lower the cost of job search effort. In contrast, a problem with non-financial subsidization (advice, encouragement, resume help, etc.) is that it might not "work" because the assistance is not that valuable or competently offered. [1] The subsidy was large enough to be meaningful but not large enough to make income effects a concern. And unlike unemployment insurance benefit extensions, which could have income effects or create incentives to delay for better matches, the only use of the subsidy is applying for jobs on the platform.

Combined with a uniquely designed subsidy, our setting and intervention have numerous advantages for understanding the behavior of new labor market entrants. Because of the computer-mediated nature of the market, we can also measure job-search behavior and outcomes essentially without error. We can observe what wage the worker proposed, who

[1] Ben Dhia, Crépon, Mbih, Paul-Delvaux, Picard and Pons (2022) shows a job search assistance website had a minimal net effect.

was hired, and what earnings, if any. We use this greater ability to observe to elucidate mechanisms that would otherwise only be conjectured. By observing earnings directly, we have a much greater ability to speak to welfare effects. And by observing the entire online market, we have a much greater ability to measure crowd-out credibly.

Our main findings are as follows. We find that directly lowering costs strongly increased job applications and, ultimately, the probability a job-seeker was hired. We find strong support for the notion that job-seekers decide how long to persist in light of the costs and perceived benefits.[2] The elasticity of applications sent to the subsidy was about 0.2. And in the most heavily subsidized group, the probability a new entrant was hired increased by nearly 30% relative to the control group.

In terms of a model of jobseeker decision-making, we show that learning and non-financial private costs matter: Workers to not simply apply until their platform-provided budget is exhausted. Many job-seekers exited far before exhausting their subsidy. The job search process is a horse race between hope and resources for new entrants. We see excess exits at points where subsidies would be exhausted. The most heavily subsidized group received 4x the initial subsidy as the control group. However, this only increases applications sent by 47%.

A 10% increase in applications is associated with about a 10% increase in the probability of being hired at least once. When we instead run this same regression but instrument for applications sent with the randomly assigned subsidy, we find that the point estimate is smaller, but not by much: a 10% increase in applications increases hire probability by about 8%. In short, there is some self-selection into greater search intensity, but it is not purely a selection effect, and job search intensity matters.

There were only tiny increases in job search intensity on the extensive margin, i.e., there was no drop-off between signing up and getting the grant and applying to a first job. This is consistent with job-seekers not knowing *ex ante* how much search effort is required to be hired. Fixed costs of starting already seem "paid"—treated workers knowing that they could apply for jobs longer (as they had many coins)—it did not seem to matter very much. This is already some evidence that workers do not internalize the full benefit of extra job search. If they did, there would be no treatment/control difference. This likely reflects that workers are initially quite uncertain.

As we can observe what jobs workers applied to, at what terms, and with what effect, we can show that not only did subsidize job-seekers apply to more jobs, they also applied to *ex ante* low success probability jobs on observable measures. In particular, more heavily subsidized job-seekers apply to more expensive jobs (by coins price), somewhat older jobs, and

[2] Somewhat surprisingly, this finding is not itself trivial or already well-known—Baker and Fradkin (2017) find no evidence that UI changes affected search behavior, as measured by Google search data, though Marinescu and Skandalis (2021) find a different result.

more jobs in different categories of work. Subsidized job seekers have a lower per-application "win rate." This reduction in win rate could also be a price effect: with the ability to send more applications, we might think that job-seekers would raise their reservation wages and bid more. However, we find no evidence this is the case—wage bids are not different from the treatment and control groups. This is consistent with workers in this setting being highly elastic, which, as we will discuss, rationalizes the lack of crowd-out.[3]

Workers receiving the highest subsidy earnings increased by 12%; had workers paid the platform's cost of subsidization, they would have still increased net earnings by 9%. This means that new entrants, on average, quit their job search too early relative to what the platform would prefer. There is little evidence of crowd-out, as determined by comparing returns to the subsidy into periods with a low and high-level of subsidy. Subsidizing the job search effort more than makes up for itself in increased job-seeker earnings. A platform taxing earnings at about 20% would find the subsidy cost-effective.

A natural question is whether these benefits were simply coming at the expense of the un-subsidized. We address crowd-out via a novel feature of the experiment. There was an initial phase where a relatively small fraction of all new market entrants were allocated to the experiment, followed by a second phase where that allocation was doubled. We find no difference in the improvement in the treatment effect on whether new entrants get started. The benefits of subsidization are entirely insensitive to this large/small allocation, undercutting the notion that crowd-out is a first-order consideration.

Despite this seemingly good news, the non-financial cost to job-seekers induced to apply more broadly almost certainly outweighed the benefit. The problem is there is still an effort cost to job applications. Job seekers still have to find jobs to apply to and send applications. As the platform is computer-mediated, we observe the time it takes to perform these tasks. Then, using the jobseekers' wage bid as a measure of their opportunity cost of time, we can calculate the implied effort cost. When we include this cost, the subsidization has a negative average welfare effect.

Ironically, jobseeker risk-aversion likely has a protective effect. Workers in this market face a difficult "cold start" problem, and buying the subsidy themselves, *ex ante*, even if it did have a positive average return, would have been a bad financial idea for most: Among the most heavily subsidized workers, only about 10% have earnings greater than the subsidy they received. We do not know if workers have good private knowledge about their prospects for success in this labor market. Still, with average benefits much larger than modal benefits, it is understandable why workers would not self-subsidize.

Our main contribution is to show the difficulty of efficiently subsidizing job search, even when the incentive can be direct. As far as we know, we are the first to report the effects of

[3]There is also some evidence they are less likely to join an "agency" which Stanton and Thomas (2016) identified as a way for new workers to overcome the cold start problem.

subsidizing job search on the intensive margin. We are the first to quantify the non-financial cost of applications and use that figure in a welfare calculation. Because of the nature of our subsidy, we can directly characterize the welfare effects of the policy, including effort costs that would otherwise be unobserved (Card, Kluve and Weber, 2010).

Given the obstacles new entrants to labor markets face, policymakers might want to support new entrants even if all gains come at the expense of incumbents (Terviö, 2009; Pallais, 2013). However, our results suggest this program, as structured, has concentrated benefits that are smaller than the diffuse costs. We contribute to a growing literature on the importance of *ex ante* beliefs in labor market job search strategies and the importance of learning (Miano, 2023; Conlon, Pilossoph, Wiswall and Zafar, 2018). It is likely that *ex ante* uncertainty about the returns to job search—combined with learning—are important in explaining variation in job search effort and outcomes. We show this in a very direct way.

The rest of the paper is organized as follows. In Section 2, we present a stylized model that highlights the difference between the platform perspective and the individual perspective on optimal job search intensity. Section 3 describes the empirical context for our study. Section 4 explores welfare. We discuss our results and conclude in Section 5.

2 Conceptual framework

We want to consider the job seeker's problem and contrast it with the social planner's problem. There are, of course, numerous models of labor market matching. We present a highly simplified and stylized model that nevertheless captures the divergence between what a social planner would choose and what a learning individual would choose in terms of effort and what an omniscient social planner would have them do.

2.1 Optimal job policy: Never start, or conditional upon starting, never quit

Suppose a risk-neutral jobseeker's probability of getting any particular job she applies to is p. Each application costs her c. All jobs are the same; she would value getting a job at v, which captures the full social benefit. Assume there is no crowd-out among job-seekers. Once she gets a job, the job search stops.

With perfect knowledge, the jobseeker applies if $pv > c$, but otherwise does not apply at all. If these values are all fixed and $p > c/v$, she keeps applying until she gets a job, which takes $1/p$ applications in expectation. This gives her an expected pay-off of $v - c/p$. The social planner would want her to follow the same rule.

As we will see empirically, many new jobseekers do not follow this rule. They quit before getting a job but still send some number of applications. One possibility is that, for some reason, they can no longer bear c, perhaps because they are credit-constrained. Another possibility is that they decide they no longer need a job because they got another job, received an inheritance, passed away, etc. As we will see, job search spells are fairly short, so these explanations seem unlikely.

Another possible explanation is that they have updated their beliefs about p, v, or c in a way so that $pv < c$. Learning about c seems unlikely, as they already know this cost from direct experience. Learning about v seems unlikely, as they have not gotten a job and have never been given a chance to know what v is "really" like. That leaves p, and this does seem like the unknown factor where learning is most likely to occur—after each rejection, the job seeker updates her beliefs about p.

The notion that job-seekers are uncertain about their *ex ante* job search returns has a strong empirical basis. Miano (2023) documents significant heterogeneity in expectations across demographic groups. For example, women expect higher costs and lower returns to job search effort. He also finds that "younger respondents believe that their search would last fewer weeks and that they will spend fewer hours searching, but they do not expect a significantly different success rate." Comparing beliefs to actual outcomes, he shows that workers have some information, but clearly, there is a large residual. Conlon et al. (2018) shows that although job-seekers are quite accurate in their assessments on average, there are substantial deviations between expectations and realized outcomes.

Returning to our stylized model, assuming she learns after each rejection, she will rationally update, lowering her p each time until she reaches a stopping condition. Although this is framed as a discrete process—and it is—you cannot send 1/2 a job application—it is useful to treat the problem continuously. It is the case that $\bar{a}$ must be an integer, the optimal decision rule is to stop when the expected value of another application equals the cost or when $\mathbf{E}[p]v = c$. In words, the job seeker stops applying when her rational belief is that any more applications would have a negative expected value. A job seeker following this strategy with an actual success probability p would have a start rate of

$$Pr(\text{Start}|\bar{a}) = 1 - (1-p)^{\bar{a}}$$

and would send, in expectation

$$\mathbf{E}[a] = \frac{1}{p}\left(1 - (1-p)^{\bar{a}}\right)$$

applications.

The difference between what an uncertain but learning jobseeker does—and what a fully-informed social planner would do—is what creates an opportunity for subsidization if the social planner has superior information. To see this difference, in Figure 1 we compare the social planner choice and individually rational choices for values of $c = 1$ and $v = 4$, with individuals with uniformly distributed p.

Figure 1: Comparison of jobseeker approaches and outcomes when per-job success probability is uniformly distributed

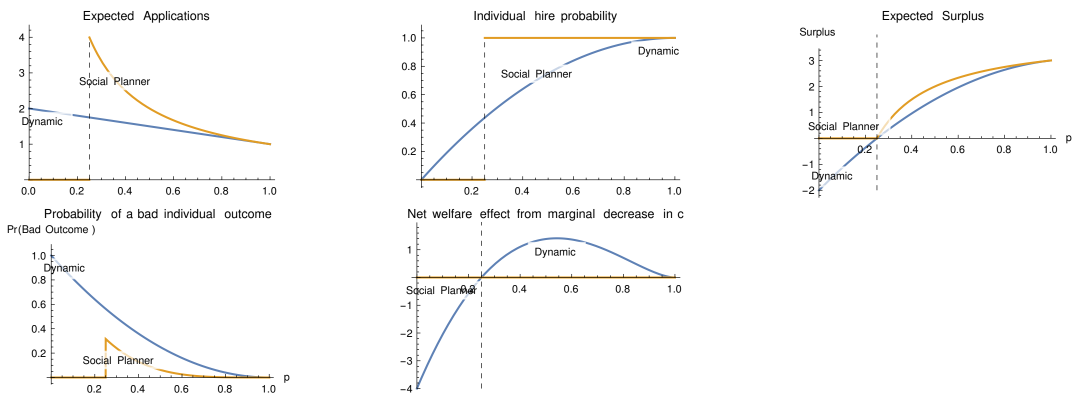

Notes: Comparison of jobseeker approaches and outcomes when per-job success probability is uniformly distributed.

Consider the number of applications sent. The learning jobseeker's number of realized applications is a declining function of p. This is because there is some probability the jobseeker gets hired before $\bar{a}$ and stops applying. In comparison, the social planner has many expected applications at the zero expectation point, with $p = c/v$ leading to v/c expected applications. The expected number is declining, reaching a value of just 1 at $p = 1$.

The effect of this difference in expected applications shows up in higher probabilities. The social planner solution has the job-seeker either never being hired or being hired for sure. By comparison, the learning jobseeker has a non-zero and increasing probability of being hired—including in ranges where the social planner would have advocated sending no applications at all.

In terms of expected surplus, the social planner is the upper envelope, reaching a max at $p = 1$ and reaching—but never going below 0 for $p < c/v$. In contrast, the individual gets a negative expected return for all $p < c/v$. And even when positive, the individual pay-off is less than the planner surplus because individual learners quit too soon.

Relatedly, we can see that while the individual is always more likely to have a bad outcome, even the social planner situation is not immune. The probability of a bad outcome is just declining in p. Only for $p = 1$ is a good outcome guaranteed.

The last thing we consider is a marginal decrease in c. For the planner, lowering c by a small amount has no effect. This is because c only enters decision-making on the extensive margin, and behavior is not changed, on the margin, by a change in c. The subsidy is helpful for the individual jobseeker when $p > c/v$; but for those with $p < c/v$, a lower c can be harmful if the p is so high they should send fewer applications rather than more. For those with $p < c/v$, lowering c makes a bad situation somewhat worse, encouraging them to keep applying when they should have had $\bar{a} = 0$. For those with $p > c/v$, lowering c helps them get closer to what they should do: set $\bar{a} = \infty$.

3 Empirical context and experiment design

The empirical context for our study is a large online labor market (Horton, 2010; Agrawal, Horton, Lacetera and Lyons, 2015; Horton, Kerr and Stanton, 2017). In online labor markets, employers hire workers to perform tasks that can be done remotely, such as computer programming, graphic design, data entry, research, and writing. Each market differs in its scope and focus, but platforms commonly provide services that include maintaining job listings, hosting user profile pages, arbitrating disputes, certifying worker skills, and maintaining feedback systems (Filippas, Horton and Zeckhauser, 2020).

Several features of conventional labor markets also exist in our context. Employers and workers are free to enter and exit the market anytime. Employers post job descriptions, and workers search and apply to jobs. Employers may invite desirable workers to apply for their jobs and can assess promising candidates through interviews. Employers and workers can negotiate over wages, which can either take the form of hourly salaries or fixed-price for projects and form contracts. More generally, employers and workers face substantial search frictions (Horton, 2017b, 2019), barriers to entry (Pallais, 2013; Stanton and Thomas, 2016), and information asymmetries (Benson, Sojourner and Umyarov, 2019; Filippas, Horton and Golden, 2021).

3.1 Starting out as a new worker

Workers are free to join the platform at any time. After verifying their identity, workers are asked to fill out their platform profiles. The platform records the degree of profile "completeness" of each worker, and displays it to prospective employers.[4]

Workers can apply directly to jobs by using up an in-platform currency called "coins." Each coin costs \$0.15, and the number of coins required to apply to a job—the cost of application—is determined by the platform using a proprietary formula. The formula only considers job-specific attributes, such as the anticipated job duration and earnings. During the experiment, workers had to spend between 1 and 6 coins to apply for a job. Employers may also invite workers to apply to jobs—although new workers seldom receive such invites—and a job application following a employer invite uses up no coins (Filippas, Horton and Pigo, 2022).

[4]Workers are aware of the factors that are used in computing the profile completeness score, which includes whether they have uploaded a profile photo, written a profile title and short description, selected skills they are proficient in, uploaded their employment history, and work experience, and so on.

Workers receive 20 coins upon joining the platform, but have to purchase coins subsequently. Workers can purchase up to 80 coins ($12) at one time. However, balances are not capped at 80 coins, and immediate repeat purchases are possible. Coins are placed in a non-interest-bearing account, cannot be converted back to cash, and expire one year after the purchase. None of the above prices or rules changed during the period of our analyses—besides the size of the initial coin grant.

3.2 Experiment

Upon registering with the platform, a fraction of job-seekers were allocated to the experiment. Workers were allocated to one of four experimental groups of that fraction. The allocation period began on September 24, 2020 and ended on December 02, 2020. A total of 753,609 workers were part of the experiment, of whom 187,635 (24.9%) were allocated to Control and received the status-quo grant of 20 coins, 189,444 (25.14%) were allocated to Treatment 1 and received a grant of 40 coins, 188,513 (25.01%) were allocated to Treatment 2 and received a grant of 60 coins, and 188,017 (24.95%) were allocated to Treatment 3 and received a grant of 80 coins. We find no evidence of systematic differences in the distributions of worker attributes by cell.

All workers in our data received the "correct" initial coin grant, according to their treatment status. The grant delivery was fast, reflected on workers' coin balance on average about 2.5 seconds after their time of allocation. Due to the fast treatment delivery, no worker in our data purchased coins before receiving the grant. Furthermore, the platform did not inform the workers that they received different grants than their peers.

3.3 Sample construction

We use both panel and cross-sectional data in our analyses. To construct the panel, we collect observations on all workers between September 24, 2020 and December 02, 2021. We then compute when each outcome was observed in days relative to the worker's allocation in the experiment, with millisecond accuracy. We divide that number by 7 and "floor" the result. As such, an application sent 6 hours after allocation takes place in period 0, an application sent 8 days after allocation takes place in period 1, and so on. We refer to period 0 as the allocation period, and we keep 52 periods for each worker. Therefore, we observe each worker's outcomes for one full year after their time of allocation.

We fill in all "missing" worker/period outcomes with 0. For example, if worker Aja sent applications in periods 1, 3, and 4, then we fill in periods 2, and 5 to 51 with 0 for the "missing" outcomes. As such, the panel is unbalanced, but the missingness is random in that it does not depend on the treatment status.

To create cross-sectional samples, we simply restrict the panel to the desired observation period and compute cumulative outcomes for each worker.

3.4 Persistence in job search effort

Figure 2 shows the histogram of total coins spent, by the group before exiting. We can see a sharp drop in spending in the control right at the 20 coins cut-off. For each active treatment, we can see excess mass at the cut-off for that group.

Figure 2: Distribution of job search effort (coins expended) by treatment group

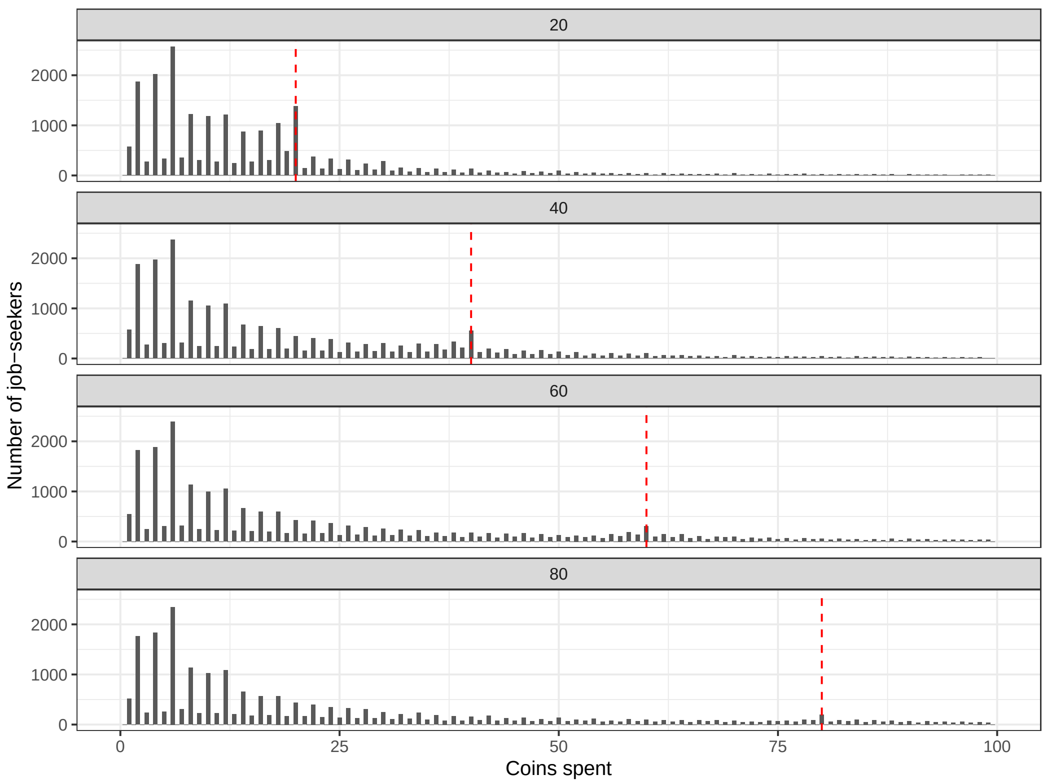

Notes: Distribution of job search effort (coins expended) by treatment group.

3.5 Panel estimates of the short-run effects of the treatment

To understand the effects of the treatment over time, we next use our data in a panel format and examine the mean differences in per-period outcomes between the experimental groups during the experimental period. Figure 3 plots estimates of the per-period differences in outcomes between each treatment group and the control group. For each period outcome, we report results as the estimated percentage change of the treatment group average over the control group average, along with a 95% confidence interval. Similarly to Section 3.7, we examine both the extensive and the intensive average treatment effect, and we winsorize the distributions of non-binary outcomes at the 99% level to deal with extreme values.

Figure 3: Panel estimates of the treatment effect on worker outcomes.

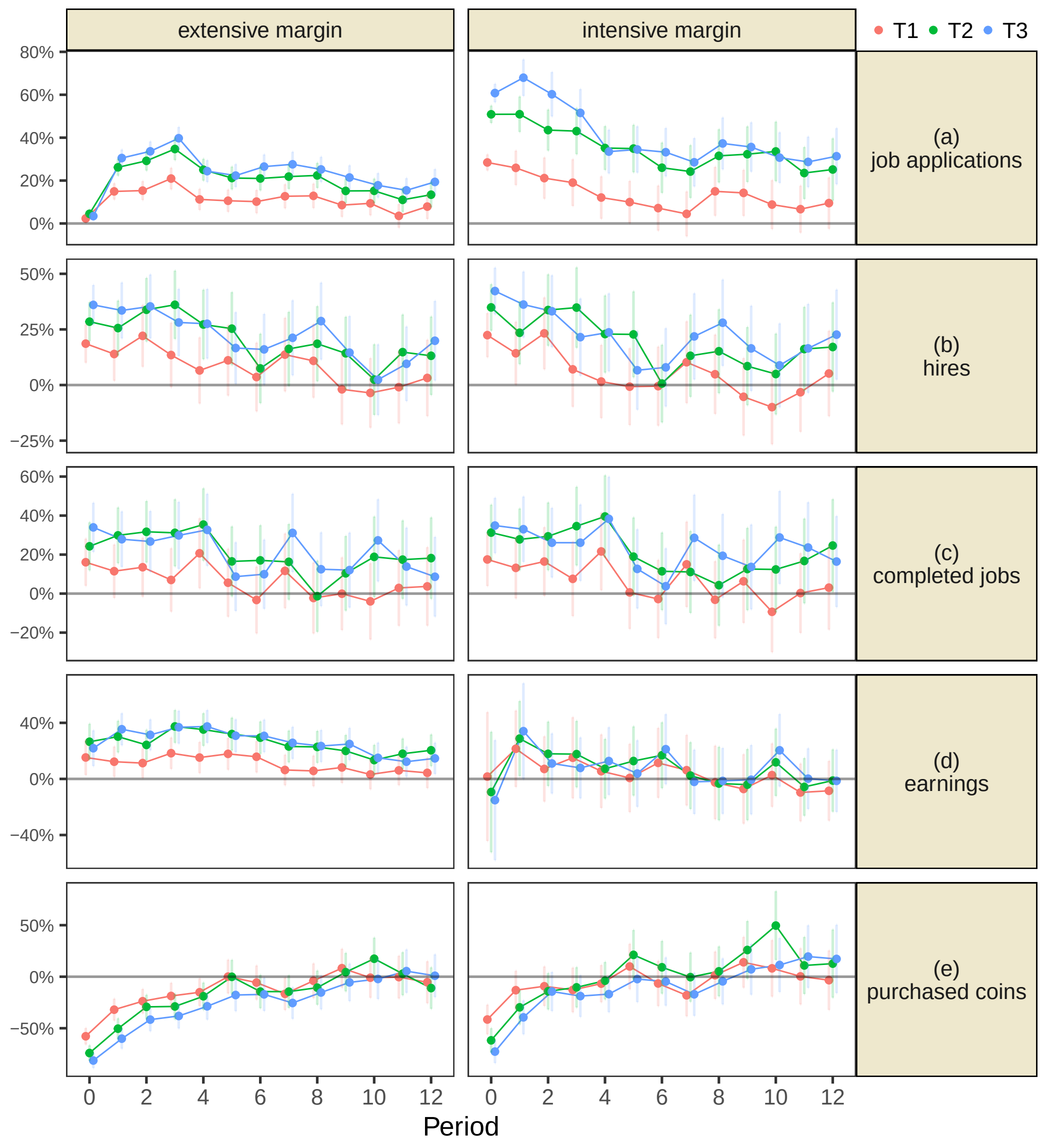

Notes: This figure plots estimates of the treatment effect on per-period worker outcomes using panel data. Each panel reports point estimates of the per-period effect of the treatment as the percentage change in the treatment group over the control group, along with a 95% confidence interval. Panels on the left examine the extensive margin effect, with the dependent variable being the indicator variable transformation of each per-period outcome. Panels on the right examine the intensive margin effect, with the dependent variable being the "raw" outcome winsorized at the 99% level. More details on the construction of the panel are provided in Section 3.3, and the results are discussed in Section 3.5.

Starting with the first period of observation (period 0), we notice that approximately the same number of workers applied to a job across all cells. However, treated workers engaged in job search much more intensely: compared to control group workers, workers in T1 placed 28.44% more bids, workers in T2 placed 50.89% more bids, and workers in T3 placed 60.79% more bids. Consequently, treated workers were substantially more likely to be hired to complete jobs and earn money early on in their platform tenure.

The effect of the treatment on workers's job search intensity persisted in subsequent periods, first increasing further and then stabilizing towards the end of the panel. In the last period of our data, the effect on job application probability compared to the control group was 7.87% for workers in T1, 13.41% for workers in T2, and 19.36% for workers in T3. Hiring outcomes followed a similar pattern to worker applications, with the effects of the treatment being positive, analogous to the initial coin grant size, and similar in magnitudes across the extensive and the intensive margins.

Treated workers were substantially less likely to buy coins in the first few periods, as they were spending down their initial grants. This effect dissipated in the later half of our panel, and workers were equally likely to buy coins after period 8 across all groups. Moreover, workers in T2 and T3 bought more coins per period by the end of the panel. Combined with the higher earnings, this result suggests that the higher initial coin grant "paid off" for the platform, at least in the short-run.

3.6 Higher subsidies cause more job applications and higher start rates

The active treatment cells persisted longer in job search, as shown in Figure 2. We can now quantify these effects on both the extensive and intensive margins. In Table 1, we regress job search effort measures on the subsidy amount.

The sample in Column (1) is all allocated job seekers. In Column (1), the outcome is whether the job-seeker sent any applications at all, on indicators for each subsidy level. The mean for the control group is just 0.163—meaning that only a bit more than 15% of users signing up apply to even one job. This is likely because the cost of signing up is low, and, after signing up, many job-seekers realize they are unlikely to be successful given the jobs listed. The fixed costs of applying do not exceed the perceived benefits.

Those starting with a larger initial coins balance applied more readily, but the effects are small. Initial application rates go from 0.163 in the control group to 0.168 in the best-performing cell, which is the 60 coins cell (not the 80 coins cell). This is about a 3% increase. While all the non-control cells are significantly better than the control, there is no apparent dose-response relationship past 20.

In Column (2), we condition upon those starting. While this is a selected sample, it is not likely too selected given the small extensive margin effects. In Column (2), among those that sent any applications, the subsidy is clearly positive, with an elasticity of about 0.17.

In Column (3) of Table 1 the outcome is whether the job-seeker eventually had on-

Table 1: Effects of subsidy on job applications sent

	Dependent variable:		
	Any applications?	Log applications	Any earnings?
	(1)	(2)	(3)
Subsidy=40	0.003*		0.002***
	(0.001)		(0.001)
Subsidy=60	0.005***		0.003***
	(0.001)		(0.001)
Subsidy=80	0.004***		0.004***
	(0.001)		(0.001)
log(Subsidy)		0.168***	
		(0.008)	
Constant	0.163***	1.138***	0.022***
	(0.001)	(0.029)	(0.0004)
All			
Apps > 0			
Observations	753,609	124,896	753,609
R^2	0.00002	0.004	0.0001

Notes: Significance indicators: $p \leq 0.1 : \ddagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$, and $p \leq .001 : ***$.

platform earnings, or "started." This is the full sample, including those that never send any applications. We can see clear evidence that the subsidy increased start rates. Furthermore, comparing across cells is a dose-response effect, with larger effects for large subsidies.

3.7 The short-run effects of the treatment

We use a cross-section that covers the first 3 months of observations for each worker (periods 0 to 12 in our panel). Throughout the rest of the paper, we refer to this period as the "experimental period."

We examine the effects of the treatment in two different ways. First, we consider the extensive margin effect treatment, applying the indicator variable transformation on each cumulative outcome. This allows us to study the fractions of workers that had non-zero outcomes during the experimental period. Second, we consider the extensive margin effect of the treatment. To deal with extreme values, we winsorize the distributions of non-binary outcomes at the 99% level.[5]

Our estimation strategy is to regress each worker outcome on indicators for the active experimental arms, i.e.,

$$y_j = \beta_0 + \beta_1 \mathrm{T}_{1j} + \beta_2 \mathrm{T}_{2j} + \beta_3 \mathrm{T}_{3j} + \epsilon,$$

[5]We report an alternative analysis of the experiment using panel data Appendix 3.5.

where y_j is the outcome of interest, T_{ij} indicates whether worker j was assigned to treatment i, and ϵ is an error term. We report the estimated effects as percentage changes over the control group outcome in Figure 4, by plotting the least squares estimates $\hat{\beta}_i/\hat{\beta}_0$ for each of the three active treatment groups, along with a 95% confidence interval around each point estimate.

Figure 4: Estimates of the effects of the treatment in the short-run.

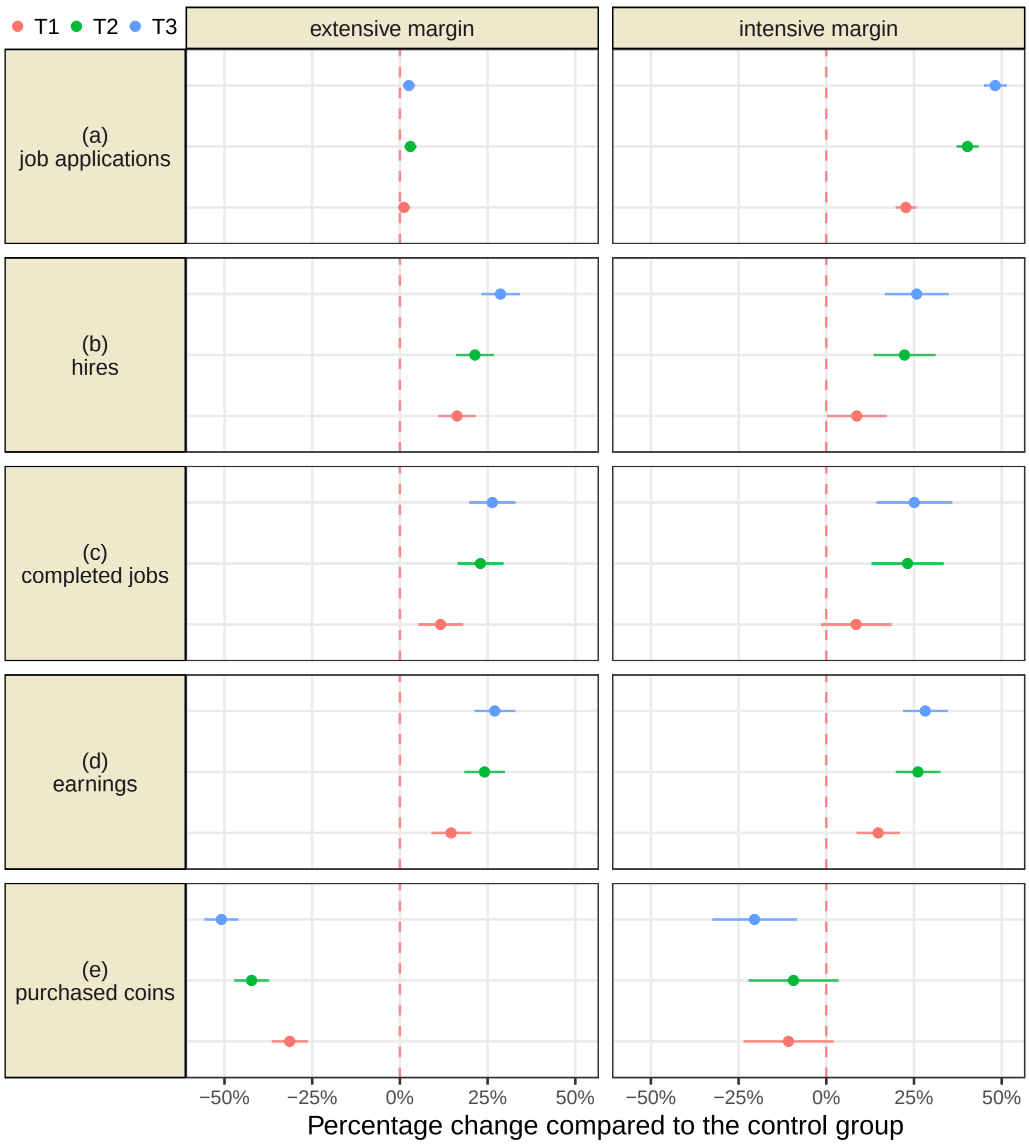

Notes: This figure plots estimates of the short-run effects of the treatment on cumulative worker outcomes, using a cross-section that covers the first 3 months of observations for each worker (periods 0 to 12). Each panel reports point estimates as the percentage change in the treatment group over the control group, along with a 95% confidence interval. Panels on the left examine the extensive margin effect, with the dependent variable being the indicator variable transformation of each outcome. Panels on the right examine the intensive margin effect, with the dependent variable being the "raw" outcome winsorized at the 99% level. See

Treated workers were more likely to apply to at least one job, but this effect was small. Compared to a baseline of 12.94% for the control group, the increase was 0.15 percentage points (1.16%) for T1, 0.39 percentage points (2.98%) for T2, and 0.33 percentage points (2.55%) for T3.[6] The estimated effects are statistically indistinguishable from one another, and conventionally statistically significant only for T2 and T3. It is worth stressing the fact that only about 1 out of 7 workers in the control group applied to at least one job and all other workers who registered never applied to a job.[7]

In contrast to the extensive margin effect, the intensive margin effect of larger grants on job applications was sizable. Compared an the average of 0.8 applications per worker in the control group, workers in T1 applied to 0.18 (22.68%) more jobs, workers in T2 applied to 0.32 (40.26%) more jobs, workers in T3 applied to 0.38 (48.16%) more jobs. These effects are statistically distinguishable from one another, and roughly analogous to the size of the initial coin grant.

Treated workers experienced better market outcomes. In particular, treated workers were more likely to be hired, to complete a job, and to start earning money during the experimental period. The effects are increasing in the size of the initial treatment, and are similar across the intensive margin. As we would expect, "moving down" the job search process, the estimated effects diminish and become more comparable across the different active treatment arms. Nevertheless, nearly all estimated effects remain conventionally statistically significant.

One negative effect for the platform was that workers receiving larger initial coin grants were substantially less likely to make a coin purchase. Compared to the control group purchase probability of 1.32%, the percentage decrease was 0.41 percentage points (31.42%) for T1, 0.41 percentage points (42.28%) for T2, and 0.67 percentage points (50.89%) for T3. However, the intensive margin estimates are smaller in magnitude than the extensive margin estimates, suggesting that some of the treated workers have already started buying coins.

3.8 The long-run effects of the treatment

The treatment had substantial and positive effects initially, but in the long run, these effects could dissipate. If only those workers who are "meant to be" remain active on the platform, and all other workers exit the platform sooner or later, then the long-run effect will be much smaller than the short-run estimates.

To examine the long-run effects of the treatment, we create a cross-sectional sample that covers the last 2 months of observation for each worker (periods 44 to 52 in the panel). We

[6]Throughout the paper, we use "percentage points" to refer to differences in levels where the outcome is naturally discussed as a fraction, and the "%" symbol to refer to percentage changes with respect to the outcome of another experimental group.

[7]This result could suggest that it is hard for new workers to start, but it could also be explained by the ease of signing up on an online labor market "on a whim."

then follow the estimation strategy of Section 3.7 to examine the extensive and intensive effects of the treatment on worker outcomes. We report the estimated effects as percentage changes in Figure 5, by plotting the least squares estimates $\hat{\beta}_i/\hat{\beta}_0$ for each of the three active treatment groups, along with a 95% confidence interval around each point estimate.

Figure 5: Estimates of the effects of the treatment in the long run

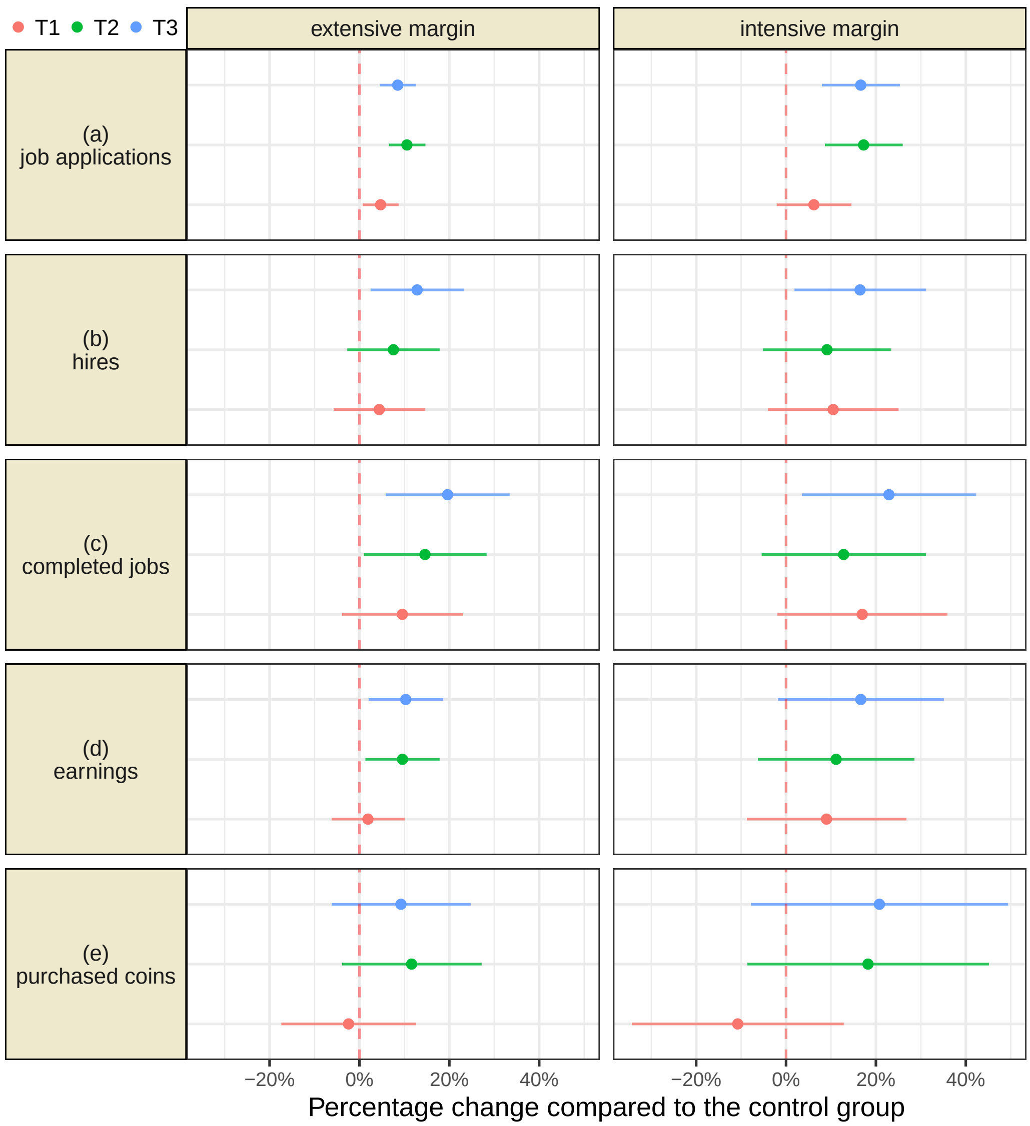

Notes: This figure plots estimates of the long-run effects of the treatment on cumulative worker outcomes, using a cross-section that covers the last 2 months of observations for each worker (periods 44 to 52). Each panel reports point estimates as the percentage change in the treatment group over the control group, along with a 95% confidence interval. Panels on the left examine the extensive margin effect, with the dependent variable being the indicator variable transformation of each outcome. Panels on the right examine the intensive margin effect, with the dependent variable being the "raw" outcome winsorized at the 99% level.

The effect of the treatment persists long after the "starting out" period. On the extensive margin, treated workers were more substantially likely to apply to a job, complete a job, or earn money. All these effect estimates are statistically significant for T3, and most of the estimates are statistically significant for T2. For T1, the estimates are positive but statistically indistinguishable from zero. On the intensive margin, the treatment seemingly also resulted in more hires, completed jobs, and earnings for all treatment groups. However, the intensive margin estimates are more imprecise, and the earnings effect estimate is not conventionally statistically significant for any treatment group. This suggests that the "raw" outcome distributions are skewed—even after winsorizing the distributions. We interpret these results as suggesting that the treatment not only attracted marginal low-earning workers but also resulted in some "heavy hitters" being retained on the platform.

4 Costs and benefits

Setting aside any concerns about crowd-out, there are at least three potential perspectives when considering the welfare of the subsidy: 1) Platform financial perspective, 2) individual financial perspective, 3) total individual perspective.

The first is straightforward: we know how much the platform spends and take in. The second is straightforward: we know how much individuals make as well. The third is hard because we need to calculate the cost of effort. But even in the first case, we need to consider that in any experiment, any benefits could come at the expense of the untreated.

In this section, we first assess crowd-out, using a unique feature of the experiment. We then consider the net welfare effects of the intervention.

4.1 Crowd-out

Crowd-out is a first-order concern in assessing efforts to subsidize job search efforts. For example, in the same platform as this field experiment, Barach et al. (2020) show that platform recommendations to employers about which of their applicants to hire clearly affect which worker gets hired, but do not affect the total number of hires. Although recommendations and increasing search efforts by job-seekers are different, this highlights the concern that experimental estimates of the effects of interventions may be highly misleading. But there is hope: also in the same empirical context, Horton (2017b) shows that encouraging employers to do more recruiting increases the probability a job is filled. Note that even if the effects of the experiment are purely due to crowd-out, new workers are so relatively disadvantaged that so long as the extra labor market income was coming at the expense of established workers, the platform may still welcome it.

Labor market crowd-out reasoning bears similarity to the "lump of labor fallacy" reasoning: There are only so many jobs, and the question is who will get them. This may be true, but the analogy to an uneconomic model of a market should give some pause. And note that this story is quite different even from the typical government fiscal policy crowd-out story, where government borrowing deleteriously affects private investment via interest rates. Fiscal policy crowd-out is a pecuniary externality. The story is *not* that the government comes in and gets the loan that was counter-factually supposed to go to some private business. This might seem like a distinction without a difference, but it matters because the extent of this crowd-out depends on the degree of elasticity, which is an empirical matter. If, for example, lenders were highly elastic in their supply of capital in keeping with the fiscal policy crowd-out story, then even extensive government borrowing might have little impact. A richer—or at least more economic—story of labor market crowd-out would be one based on pecuniary externalities and supply elasticities.

As in all labor markets, workers likely do not fully internalize the benefits of job search effort, or only do as a knife-edge condition (Diamond, 1982; Hosios, 1990; Mortensen and Pissarides, 1994) Encouraging higher job search intensity from select workers can be in the platform's interest, as the platform taxes the wage bill from formed matches. Of course, there are downsides to encouraging job-search, such as crowd-out (Crépon et al., 2013; Horton and Vasserman, 2020)—over-application is considered enough of a problem that the platform has a price for applying to jobs. However, whether job search behavior *should* be changed is a separate question from whether job search behavior *can* be changed.

To measure the extent of crowd-out, we can exploit a feature of the experiment. Figure 10 shows the daily allocations of experimental subjects to groups over time. Note that the experiment was conducted in two phases that differed dramatically in the size of subjects allocated.

As we saw earlier, the higher subsidies strongly increase start rates, but there is no evidence that the treatment differed in effectiveness by time period. If the treatment worked primarily by taking hires from those not treated, we should see a reduction in treatment effects in the second period. That is empirically not the case. If we nest one model in the other and do a likelihood ratio test, the p-value is 0.06.

4.2 Welfare

For each subsidy level, we can compute the increase in earnings from the assigned workers compared to the control. It is positive in all cells. The effects are especially large in T2, but recall that these are cumulative effects. The cost of the subsidy is computed by calculating the net difference in coins sales. The ratio of the two is the return on investment. Table 3 computes the total increase in earnings, the total cost of subsidization, and the implied return on investment.

Table 2: Effects of treatment over time

	Dependent variable:	
	Any hires?	
	(1)	(2)
Subsidy=20	0.013***	0.013***
	(0.0003)	(0.0003)
Subsidy=40	0.002***	0.002***
	(0.0005)	(0.0004)
Subsidy=60	0.004***	0.003***
	(0.0005)	(0.0004)
Subsidy=80	0.004***	0.004***
	(0.0005)	(0.0004)
Early Period	0.001*	0.001***
	(0.001)	(0.0003)
Subsidy=40 x Early	0.001	
	(0.001)	
Subsidy=60 x Early	−0.001	
	(0.001)	
Subsidy=80 x Early	0.001	
	(0.001)	
Observations	753,609	753,609
R^2	0.0002	0.0002

Notes: This table reports regressions where the dependent variables Significance indicators: $p \leq 0.1$: ‡, $p \leq 0.05 : *$, $p \leq 0.01 : **$, and $p \leq .001 : ***$.

Table 3: Efficiency of subsidization

Experimental group	Increase in Earnings	Subsidy Cost	ROI	BE Tax	BE per-app Hassle Cost
Treatment1	\$22,621.58	\$3,812.85	4.93	16.85%	\$0.45
Treatment2	\$73,302.47	\$3,332.10	21.00	4.55%	\$0.79
Treatment3	\$68,276.66	\$7,524.60	8.07	11.02%	\$0.58

While a social planner would not care, the platform only captures some of the revenue through "taxation." We compute the tax rate on labor income that would make this subsidization self-financing. This would lead to decision-making similar to a budget-constrained government. We can see that tax rates as low as 5% make subsidization worth it. If the planner/platform put extra weight on new workers starting, this implied tax rate would be lower.

This platform perspective does not consider the marginal cost of application to the jobseeker. However, as we know how many more applications were stimulated by the subsidy, we can calculate what per-application hassle cost would be needed to dissipate the benefit of the subsidy. This number is reported in the far right column of the table.

4.3 Accounting for jobseeker effort costs

We know that the subsidy appeared to have a positive return on investment for the platform. And average financial pay-off was higher for treated workers—though some were induced to send more applications and were not better off. Furthermore, this perspective does not consider the hassle cost of a job application. In this section, we consider these costs.

We do not observe the non-financial costs of job search directly. Although job search has psychological and stress-related costs, the primary marginal cost is time. Each additional application requires finding that job, reading the job description, preparing a cover letter, and ultimately applying. Our marketplace is computer-mediated, which means we can measure the time an application was sent down to the millisecond. If we look at the elapsed time between subsequent job posts, that might be informative about the time it takes to find and apply for another job.

Of course, we need to be careful about finding true "spells" of job search: If we find that a jobseeker applied on one day and then applied again three days later, we should not presume she spent 36 hours crafting her application. To deal with this issue, we only look at job applications that happened on the same day and plot the distribution. We take all jobseekers that sent more than two applications on the same day. We then compute the elapsed time between applications, down to the second.

Here is a sample of the data for some worker with ID = 1. Note that application 2 comes more than 20 hours after application 1. But application 4 is just a little more than 30 minutes later:

```
           ID bid_ts              bid_date   hourly_rate num.obs delta.t
        <int> <dttm>              <IDate>          <dbl>   <int>   <dbl>
1           1 2020-09-24 02:46:56 2020-09-24          35       3 NA
2           1 2020-09-24 23:31:20 2020-09-24          40       3 20.7
3           1 2020-09-24 23:52:49 2020-09-24          40       3  0.358
4           1 2020-09-25 21:48:23 2020-09-25          50       3 NA
```

Note that we also have an hourly bid, which we can use to value the bid.

If we do this for a large number of workers, we can get a distribution of elapsed times between applications. This distrbution is plotted on a log scale in Figure 6.

Figure 6: Distribution of elapsed time between job applications

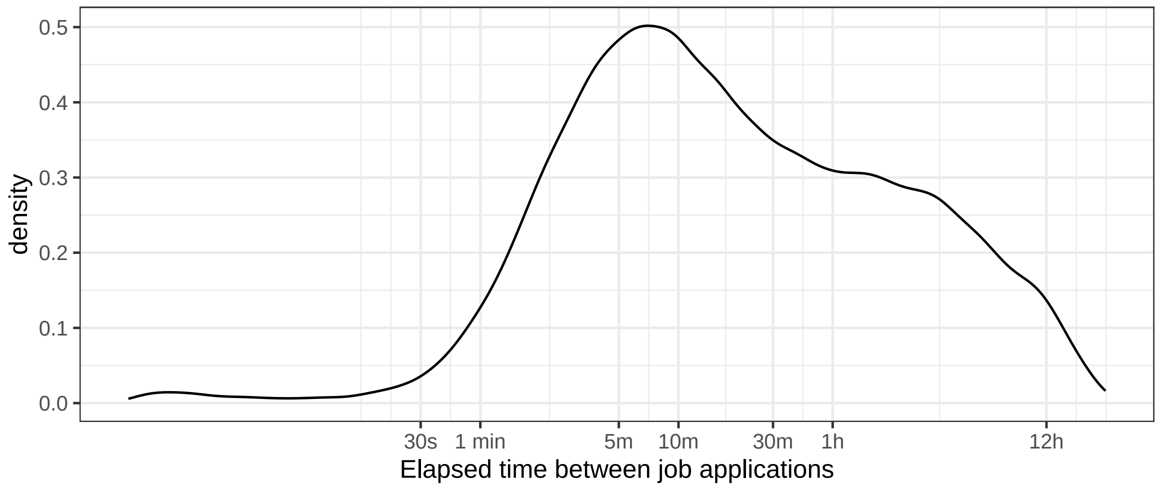

Notes: This plot shows the elapsed time between job applications.

We can see there is a uni-modal distribution centered around 7.61 minutes. Notice that the left part of the distribution looks smooth while the right tail is much less so. It seems likely that this is a mixture distribution of the true time it takes an applicant to prepare an application (which is approximately log-normally distributed) as a distribution of times that reflect applications that span "spells," e.g., the user sent some applications in the morning and then again in the evening.

If we can separate out these "spell breaks" we can estimate the distribution of application preparation times. But once we have this time measurement, how do we value it? Fortunately, each application to an hourly job has an associated wage bid. While this is aspirational, it is not likely too far away from a worker's reservation wage in a competetive market.

To deal with the mixture distirbution problem, we estimate the parameters of a log normal distribution but only using the the part of the distribution that likely does not reflect picking up "spells." We do this by setting a t_{max} (e.g., 10 minutes) and assuming that all t's less than this are from the true distribtion of application costs, while those to the right are a mixture of observations of application writing costs and breaks. While this is imperfect—presumably there are real breaks of time less than t_{max} within applications, presumably most are to the right.

For this reason, we focus on the median value of the distribution, which more likely represents the mean application cost (in time) faced by a jobseeker.

To convert from a measure of time to money, we need to use the jobseekers time value of money.

In this data, the median wage bid is \$15/hour. Surprisingly, there is no discernible relationship between application time and wage bid.

Figure 7: Distribution of elapsed time between job applications

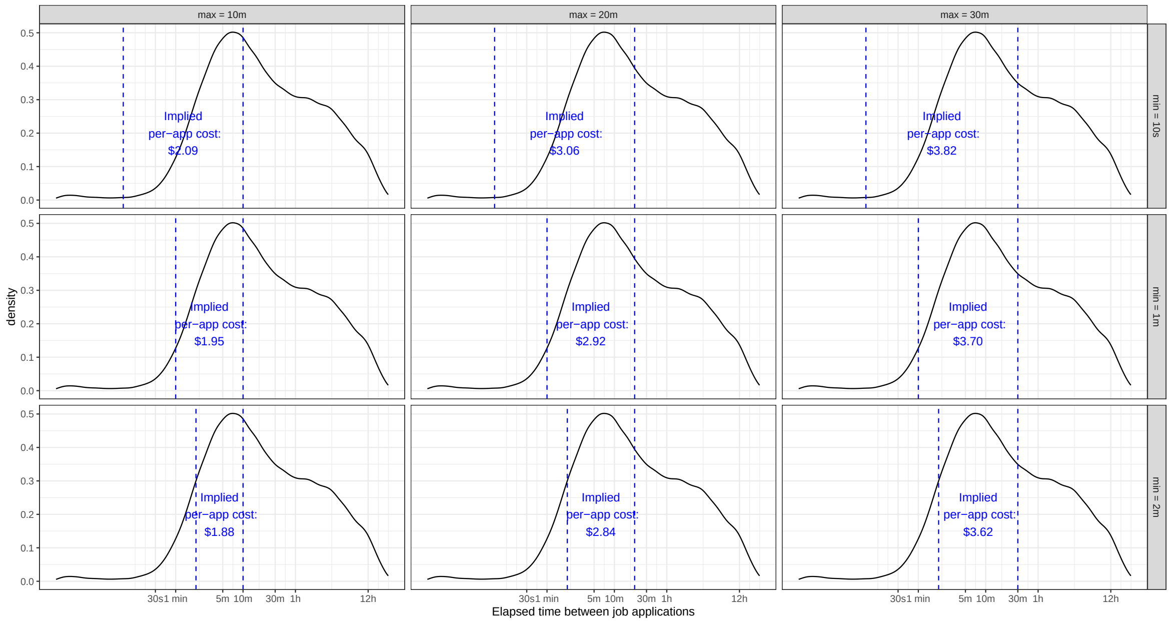

Notes: This plot shows the elapsed time between job applications.

If we take the median wage as the opportunity cost of time, then with \$15/hour and 7.61 minutes, the hassle cost of an application is \$1.9. This is considerably larger than the break-even per-application cost in Table 3, suggesting that the expected value of the subsidy from a planner perspective is negative. The hassle cost from the increase in applications likely outweighed the expected financial benefits.

5 Conclusion

This paper reports the results of a field experiment that subsidized the job search of new labor market entrants. The main effect of the subsidy was is to increase the number of job applications sent, which in turn increased earnings. Despite an increase in earnings that greatly outweighed the direct financial cost of the subsidy, the increase in non-financial application costs likely dissipated the benefits.

Despite this disappointing result, the desirability of this kind of subsidy depends in part on the benefit attributed to new job seekers getting started. If some of the benefits come at the expense of established workers, this is likely attractive from an equity standpoint, as new labor market entrants are particularly disadvantaged.

If there was a desire to mirror this program in the broader labor market, an obvious complication is that the subsidy in our setting is the reduction of a per-application financial cost, which is superficially unlike subsidization in conventional settings. In conventional settings, there is typically no direct financial cost to applying for a job, only the hassle or ordeal cost. However, UI payments contingent on proven job search effort is commonplace in many US state systems, directly subsidizing job search effort.[8] For administrative and monitoring reasons, these schemes are structured as benefits being contingent upon applying to some number of jobs, not a per-job cash payment.[9] When payments are not contingent upon job search effort, the idea or hope is that benefits are not so generous as to make the labor/leisure trade-off such that the job seeker does not apply for jobs—nor so stingy that job-seekers form bad matches. Our subsidization is attractive because using a platform currency that can only be used to apply to jobs eliminates the fraud concern but still creates incentives on the intensive margin. One could imagine a baseline UI payment for no effort, then a schedule of increasing generosity for more jobs applied to. The increasing computer mediation of all aspects of the labor market would, in principle, make more complex ways of structuring UI payments possible and more like the system described in this paper.

[8] https://www.mass.gov/service-details/work-search-examples

[9] A cash payment per job applied would likely lead to fraud, though if it could be structured appropriately, creating an intensive margin would create good incentives. https://www.businessinsider.com/mcdonalds-pays-50-for-job-interviews-highlighting-hiring-struggles-2021-4

References

Agrawal, Ajay, John J Horton, Nicola Lacetera, and Elizabeth Lyons, "Digitization and the contract labor market: A research agenda," in "Economic analysis of the digital economy," University of Chicago Press, 2015, pp. 219–250.

Baker, Scott R. and Andrey Fradkin, "The Impact of Unemployment Insurance on Job Search: Evidence from Google Search Data," *The Review of Economics and Statistics*, 2017, *99* (5), 756–768.

Barach, Moshe A., Joseph M. Golden, and John J. Horton, "Steering in Online Markets: The Role of Platform Incentives and Credibility," *Management Science*, 2020, *66* (9), 4047–4070.

Beam, Emily A., "Search costs and the determinants of job search," *Labour Economics*, 2021, *69*, 101968.

Benson, Alan, Aaron Sojourner, and Akhmed Umyarov, "Can reputation discipline the gig economy? Experimental evidence from an online labor market," *Management Science*, 2019.

Card, David, Jochen Kluve, and Andrea Weber, "Active labour market policy evaluations: A meta-analysis," *The Economic Journal*, 2010, *120* (548), F452–F477.

Conlon, J. J., L. Pilossoph, M. Wiswall, and B. Zafar, "Labor market search with imperfect information and learning," Technical Report, National Bureau of Economic Research 2018.

Crépon, Bruno, Esther Duflo, Marc Gurgand, Roland Rathelot, and Philippe Zamora, "Do labor market policies have displacement effects? Evidence from a clustered randomized experiment," *The Quarterly Journal of Economics*, 2013, *128* (2), 531–580.

DellaVigna, Stefano, Jörg Heining, Johannes F Schmieder, and Simon Trenkle, "Evidence on Job Search Models from a Survey of Unemployed Workers in Germany*," *The Quarterly Journal of Economics*, 10 2021, *137* (2), 1181–1232.

Dhia, A. Ben, B. Crépon, E. Mbih, L. Paul-Delvaux, B. Picard, and V. Pons, "Can a website bring unemployment down? experimental evidence from france," Technical Report, National Bureau of Economic Research 2022.

Diamond, Peter A, "Aggregate demand management in search equilibrium," *Journal of political Economy*, 1982, *90* (5), 881–894.

Filippas, Apostolos, John Horton, and Zmetro Pigo, "Advertising as Coordination," *Working paper*, 2022, *0* (0), 000–00.

_ **, John J Horton, and Joseph Golden**, "Reputation Inflation," *Marketing Science*, 2021.

_ **,** _ **, and Richard J Zeckhauser**, "Owning, using, and renting: Some simple economics of the "sharing economy"," *Management Science*, 2020, *66* (9), 4152–4172.

Horton, John J, "Online labor markets," *Internet and Network Economics*, 2010, pp. 515–522.

_ , "The Effects of Algorithmic Labor Market Recommendations: Evidence from a Field Experiment," *Journal of Labor Economics*, 2017, *35* (2), 345–385.

_ , "The effects of algorithmic labor market recommendations: Evidence from a field experiment," *Journal of Labor Economics*, 2017, *35* (2), 345–385.

_ , "Buyer uncertainty about seller capacity: Causes, consequences, and a partial solution," *Management Science*, 2019.

_ **and Shoshana Vasserman**, "Job-Seekers Send Too Many Applications," *Working paper*, 2020.

_ **, William R Kerr, and Christopher Stanton**, "Digital labor markets and global talent flows," Technical Report, National Bureau of Economic Research 2017.

Hosios, Arthur J, "On the efficiency of matching and related models of search and unemployment," *The Review of Economic Studies*, 1990, *57* (2), 279–298.

Jäger, S., C. Roth, N. Roussille, and B. Schoefer, "Worker beliefs about outside options," Technical Report, National Bureau of Economic Research 2021.

Mangin, Sephorah and Benoît Julien, "Efficiency in search and matching models: A generalized Hosios condition," *Journal of Economic Theory*, 2021, *193*, 105208.

Marinescu, Ioana and Daphné Skandalis, "Unemployment insurance and job search behavior," *The Quarterly Journal of Economics*, 2021, *136* (2), 887–931.

Miano, Armando, "Search Costs, Outside Options, and On-the-Job Search," *Working paper*, 2023.

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*, 2013.

Stanton, Christopher T and Catherine Thomas, "Landing the first job: The value of intermediaries in online hiring," *The Review of Economic Studies*, 2016, *83* (2), 810–854.

Terviö, Marko, "Superstars and Mediocrities: Market Failure in the Discovery of Talent1," *The Review of Economic Studies*, 04 2009, *76* (2), 829–850.

A Additional experimental details and results

A.1 The randomized assignment was performed correctly

Figure 10 shows the daily allocations of workers to experimental groups over time. Note that the experiment was conducted in two phases that differed dramatically in the percentage of all new workers who were allocated to the experiment.

Figure 8: Workers allocated to the control and treatment groups over time

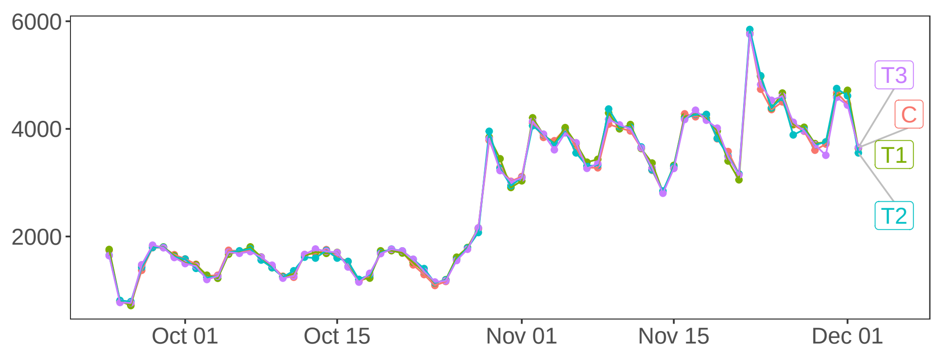

Notes: This figure plots the number of workers allocated to the control and treatment groups each day of the allocation period. The allocation period began on September 24, 2020 and ended on December 02, 2020.

Table 5 reports two-sided t-tests for various worker-level attributes for one active treatment arm and the control group. We find no evidence of systematic differences between the experimental groups.

Table 5: Balance test table

	Control mean $\overline{X}_{CTL}$	Treatment 3 mean $\overline{X}_{T3}$	p-value
Worker Characteristics			
email length	23.23	23.23	0.819
yahoo email	0.03	0.03	0.776
google email	0.83	0.83	0.493
US-based	0.19	0.19	0.51
UK-based	0.03	0.03	0.091
Observation counts	187,635	188,017	0.533

Notes: This table reports averages and p-values of two-sided t-tests for various pre-treatment observables, for workers assigned to the control group and to treatment T3. The reported attributes are, for each worker, (i) the number of characters in their email, (ii) whether their email is a Yahoo email, (iii) whether their email is a Google email, (iv) whether they are based in the USA, and (v) whether they are based in the UK. Performing the same tests for other experimental arms yields no evidence of systematic differences.

B Workers' job search strategies

We examine next how the treatment affected workers' job search strategies. To that end, we focus on workers who applied to at least one job during the experimental period. Because of this restriction, our approach is necessarily observational. However, it is worth noting that the effect of the treatment on worker entry was very small (see Section **??**). Furthermore, we find no evidence of systematic differences in observable pre-treatment characteristics between workers who applied to at least one job during the experimental period, and who were assigned to different experimental cells (see Appendix **??**).

B.1 First job search spell

We begin by examining aggregate outcomes related to the workers' first job search spell. We define the first job search spell to begin at the time of a worker's first job application, and to end either at the time of the worker' last application or at the time of the worker's first job hire—whichever comes first during the experimental period. Our estimation strategy is to regress each outcome on indicators for the active experimental arms. To deal with extreme values, we winsorize the distributions of non-binary outcomes at the 99% level. Figure 9 reports the estimated effects for each of the three active treatment groups as percentage changes over the control group outcomes. We report all regression tables in Appendix **??**.

Treated workers performed both a more intensive and a more extensive job search. On the intensive side, the duration of treated workers' first job search spell was 6.68% more days for T1, 9.84% more days for T2, and 11.88% more days for T3, compared to the control group. During that spell, treated workers applied to substantially more jobs and used more coins, with the size of this effect being analogous to the size of the initial coin endowment.[10] On the extensive side, treated workers applied to a more diverse group of job categories. Compared to the control group, the estimated increase is 9.32% for T1, 15.14% for T2, and 18.21% for T3. Similarly, treated workers were more likely to apply to jobs requiring a wider range of coin fees, with the effect being again analogous to the size of the initial coin endowment. As a result, treated workers were much more likely to be hired by the end of their first job search spell.

[10]It is worth noting that the increase in the number of applications is similar to that estimated during the entire experimental period

Figure 9: Initial coin grant and the first job search spell (I)

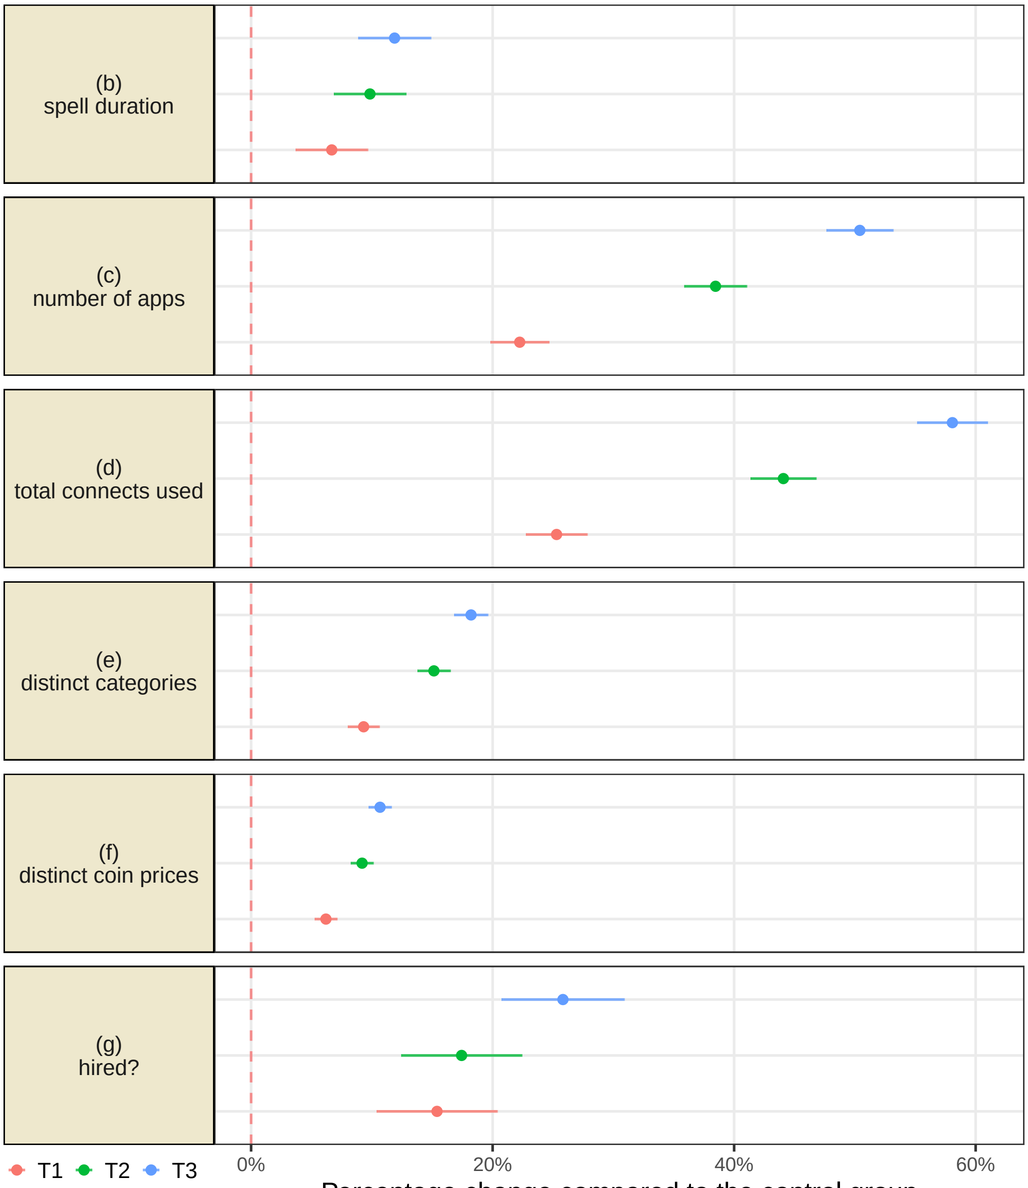

Notes: This figure plots estimates of the treatment effect on workers' first job search spell outcomes. The sample comprises workers who applied to least one job during the experimental period, and the distributions of non-binary dependent variables are winsorized at the 99% level. Each panel reports point estimates of the effect of the treatment as the percentage change in the treatment group over the control group, along with a 95% confidence interval. More details on the construction of the sample, dependent variable definitions, and a discussion of the results are provided in Section B.1. All regression tables in Appendix ??.

C Regression tables

Table 6: Cross-sectional estimates of the extensive margin effects of the treatment on workers's outcomes in the short-run.

	Dependent variable:				
	applications	hires	completed jobs	earnings	purchased coins
	(1)	(2)	(3)	(4)	(5)
Control	0.129***	0.015***	0.011***	0.014***	0.013***
	(0.001)	(0.0003)	(0.0002)	(0.0003)	(0.0003)
Treatment 1	0.002	0.003***	0.001***	0.002***	−0.004***
	(0.001)	(0.0004)	(0.0003)	(0.0004)	(0.0003)
Treatment 2	0.004***	0.003***	0.002***	0.003***	−0.006***
	(0.001)	(0.0004)	(0.0004)	(0.0004)	(0.0003)
Treatment 3	0.003**	0.004***	0.003***	0.004***	−0.007***
	(0.001)	(0.0004)	(0.0004)	(0.0004)	(0.0003)
Observations	753,609	753,609	753,609	753,609	753,609
R^2	0.00002	0.0001	0.0001	0.0001	0.001

Notes: This table reports regressions where the dependent variables are indicator-variable-transformed worker cumulative outcomes in the cross-sectional data. The independent variables are treatment indicators. The reported outcomes are workers' (1) job applications, (2) job hires, (3) completed jobs, (4) earnings, and (5) purchased coins. The left-hand-side panels of Figure 4 plot the treatment effects for the treatment cells as percentage changes over the control group estimate. Significance indicators: $p \leq 0.1 : \ddagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$, and $p \leq .001 : ***$.

Table 7: Cross-sectional estimates of the intensive margin effects of the treatment on workers's outcomes in the short-run.

	Dependent variable:				
	applications	hires	completed jobs	earnings	purchased coins
	(1)	(2)	(3)	(4)	(5)
Control	0.799***	0.036***	0.024***	0.529***	1.309***
	(0.008)	(0.001)	(0.001)	(0.011)	(0.055)
Treatment 1	0.181***	0.003*	0.002‡	0.078***	−0.141‡
	(0.012)	(0.002)	(0.001)	(0.017)	(0.085)
Treatment 2	0.322***	0.008***	0.006***	0.138***	−0.122
	(0.013)	(0.002)	(0.001)	(0.017)	(0.085)
Treatment 3	0.385***	0.009***	0.006***	0.149***	−0.268***
	(0.013)	(0.002)	(0.001)	(0.017)	(0.080)
Observations	753,609	753,609	753,609	753,609	753,609
R^2	0.001	0.0001	0.00004	0.0001	0.00001

Notes: This table reports regressions where the dependent variables are worker cumulative outcomes in the cross-sectional data, winsorized at the 99% level. The independent variables are treatment indicators. The reported outcomes are workers' (1) job applications, (2) accepted job applications, (3) completed jobs, (4) earnings, and (5) purchased coins. The right-hand-side panels of Figure 4 plot the treatment effects for the treatment cells as percentage changes over the control group estimate. Significance indicators: $p \leq 0.1 : \ddagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$, and $p \leq .001 : * * *$.

Table 8: Cross-sectional estimates of the extensive margin effects of the treatment on workers's outcomes in the long-run.

	Dependent variable:				
	applications	hires	completed jobs	earnings	worker activity
	(1)	(2)	(3)	(4)	(5)
Control	0.026***	0.004***	0.002***	0.006***	0.027***
	(0.0004)	(0.0001)	(0.0001)	(0.0002)	(0.0004)
Treatment 1	0.001*	0.0002	0.0002	0.0001	0.001*
	(0.001)	(0.0002)	(0.0002)	(0.0003)	(0.001)
Treatment 2	0.003***	0.0003	0.0003*	0.001*	0.003***
	(0.001)	(0.0002)	(0.0002)	(0.0003)	(0.001)
Treatment 3	0.002***	0.001*	0.0005**	0.001*	0.002***
	(0.001)	(0.0002)	(0.0002)	(0.0003)	(0.001)
Observations	753,609	753,609	753,609	753,609	753,609
R^2	0.00004	0.00001	0.00001	0.00001	0.00004

Notes: This table reports regressions where the dependent variables are indicator-variable-transformed worker cumulative outcomes in the cross-sectional data. The independent variables are treatment indicators. The reported outcomes are workers' (1) job applications, (2) job hires, (3) completed jobs, (4) earnings, and (5) whether the worker was active during the cross-section. The left-hand-side panels of Figure 5 plot the treatment effects for the treatment cells as percentage changes over the control group estimate. Significance indicators: $p \leq 0.1 : \ddagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$, and $p \leq .001 : ***$.

Table 9: Cross-sectional estimates of the intensive margin effects of the treatment on workers's outcomes in the long-run.

	Dependent variable:			
	applications	hires	completed jobs	earnings
	(1)	(2)	(3)	(4)
Control	0.314***	0.009***	0.005***	6.956***
	(0.009)	(0.0005)	(0.0003)	(0.446)
Treatment 1	0.019	0.001	0.001‡	0.626
	(0.013)	(0.001)	(0.0005)	(0.628)
Treatment 2	0.054***	0.001	0.001	0.774
	(0.014)	(0.001)	(0.0004)	(0.615)
Treatment 3	0.052***	0.002*	0.001*	1.157‡
	(0.014)	(0.001)	(0.0005)	(0.652)
Observations	753,609	753,609	753,609	753,609
R^2	0.00003	0.00001	0.00001	0.00000

Notes: This table reports regressions where the dependent variables are worker cumulative outcomes in the cross-sectional data, winsorized at the 99% level. The independent variables are treatment indicators. The reported outcomes are workers' (1) job applications, (2) accepted job applications, (3) completed jobs, and (4) earnings. The right-hand-side panels of Figure 5 plot the treatment effects for the treatment cells as percentage changes over the control group estimate. Significance indicators: $p \leq 0.1 : \ddagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$, and $p \leq .001 : ***$.

Table 10: Cross-sectional estimates for workers' first job search spell strategies.

	Dependent variable:					
	duration	total apps	coins used	categories	app prices	hired?
	(1)	(2)	(3)	(4)	(5)	(6)
Control	14.231***	5.503***	18.198***	2.204***	2.243***	0.118***
	(0.153)	(0.045)	(0.157)	(0.010)	(0.007)	(0.002)
Treatment 1	0.950***	1.224***	4.604***	0.205***	0.139***	0.018***
	(0.217)	(0.068)	(0.236)	(0.015)	(0.011)	(0.003)
Treatment 2	1.401***	2.117***	8.022***	0.334***	0.206***	0.021***
	(0.217)	(0.073)	(0.252)	(0.015)	(0.011)	(0.003)
Treatment 3	1.691***	2.774***	10.569***	0.402***	0.240***	0.030***
	(0.218)	(0.078)	(0.270)	(0.016)	(0.011)	(0.003)
Observations	99,073	99,073	99,073	99,073	99,073	99,073
R^2	0.001	0.014	0.017	0.007	0.006	0.001

Notes: This table reports regressions where the dependent variables are worker outcomes during their first job search spell. The independent variables are treatment indicators. The sample consists of workers who placed at least one bid during the experimental period, and is restricted to the time between allocation to the treatment, and the minimum of the first hire (if any), or otherwise the end of the experimental period. For each worker's first job search spell, the reported outcomes are (1) the duration, (2) the total number of job applications, (3) the total number of coins expended, (4) the number of distinct categories of jobs applied to, (5) the number of distinct application prices, and (6) whether the first job search spell ended in a hire. Figure 9 plots the treatment effects for the treatment cells as percentage changes over the control group estimate. Significance indicators: $p \leq 0.1 : \ddagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$, and $p \leq .001 : ***$.

D Misc analyses

D.1 Poisson regresisons

Figure 10: Poisson regression of key job-seeker outcomes

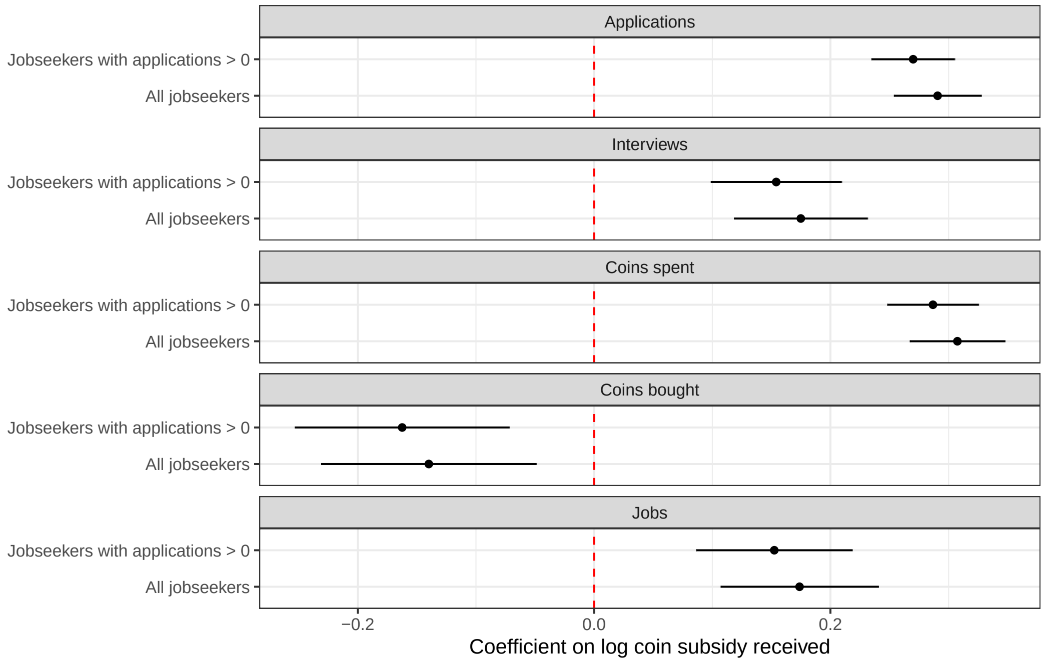

Notes:

D.2 How getting a job affects wage bids?

Table 11: How getting a job affects wage bids

	Wage bid (Log) (1)
After First Hire	0.0567***
	(0.0040)
R^2	0.839
Observations	1,605,682
Freelancer fixed effects	✓

Notes: Significance indicators: $p \leq 0.1 : \ddagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$, and $p \leq .001 : ***$.