LIQUIDITY FOR TEACHERS: EVIDENCE FROM TEACH FOR AMERICA AND LINKEDIN*

LUCAS C. COFFMAN
JOHN J. CONLON
CLAYTON R. FEATHERSTONE
JUDD B. KESSLER
JESSICA MIXON

March 6, 2023

There are teacher shortages in the U.S. and around the world. In a three-year field experiment with a large teacher placement program, Teach For America (TFA), Coffman, Conlon, Featherstone and Kessler (2019) finds that providing upfront liquidity to prospective teachers in financial need dramatically increases the rate at which they start teaching through TFA. In this paper, we combine TFA administrative data, survey data, and publicly available data (e.g., LinkedIn profiles) to extend those results. We follow individuals for a few years post treatment and find that providing upfront liquidity not only increases the rate that financially constrained individuals join TFA but also increases the rate that they complete the full two years of teaching. Further, providing liquidity to those who need it increases their likelihood of being teachers at all—not just through TFA—through at least two years.

JEL Codes: I21, J22, J45, J62, J68.

*The authors thank their partners at Teach For America: Demi Ross, Sean Waldheim, Alex Spangler, Lee Loyd, Johann von Hoffmann, Lauren Moxey, and Brigid Peña. This project was only possible via funding from the Wharton Behavioral Lab and the Center for Human Resources at Wharton. Helpful feedback was provided by Alexandra Opanasets and Ryan M Westphal.

I. INTRODUCTION

The United States is currently facing a teacher shortage. Over half of U.S. states reported a shortage of teachers in the 2018–2019 school year (Learning Policy Institute, 2019). Recent data from the U.S. Department of Education suggests that the number of people entering the teaching profession has decreased, with enrollment in teacher preparation programs dropping by 35% between 2009 and 2014 (US Dept Ed, 2018). This shortage has led to larger classrooms, an increase in the number of teachers working outside of their subject area of expertise and an increase in the number of teachers with emergency or provisional credentials. Further, the shortage is even more severe for teachers of color and teachers from low-socioeconomic backgrounds. As a result, African American, Hispanic, and American Indian students are more likely to have teachers of a different race than are White students (Learning Policy Institute, 2018b). As shown in Dee (2004), having a race-matched teacher improves student performance.[1] Additionally, low-income students are more likely to be taught by under-prepared, out-of-field, and inexperienced teachers (Learning Policy Institute, 2018a).

Recent experimental evidence in Coffman et al. (2019) identified a promising, implementable policy to attract more teachers: Providing upfront liquidity to potential teachers. If some potential teachers are liquidity constrained, providing modest cash-on-hand before various costs are incurred can allow them to bridge the gap until they receive their first paycheck. This paper follows up on the prospective teachers from that experiment. Whereas the original

[1]Though we cannot report the race of teachers in the experiment described, the treatment effects found are for lower-SES teachers, which is correlated with many factors including race.

work showed effects on teaching on the first day of the first year of teaching, we aim to understand how providing upfront liquidity can increase the number of teachers one, two, or three years after the liquidity provision. Additionally, we measure the effect not just for teaching through the organization providing the liquidity, but on the overall number of teachers.

Coffman et al. (2019) reports on a three-year field experiment with Teach For America (TFA), a large teacher-placement, training, and support program in the United States. Approximately half of the prospective teachers admitted to TFA apply for funding through a "Transitional Grants and Loans" (TGL) program run by TFA, which financially supports TFA teachers' transitions into teaching. The experiment described in that paper introduces random variation in the funding package offered to potential teachers and observes whether they join TFA and begin teaching. While the majority of the TGL applicants are unaffected by marginal increases in the funds offered to them, those with the highest financial need are substantially more likely to join TFA if offered even a few hundred dollars more by the program. The paper finds that additional grant and additional loan offers are equally effective at encouraging individuals to join TFA and reports on survey data in which those with the highest financial need report limited access to credit markets. The paper concludes that the funds induce individuals to become teachers because they loosen liquidity constraints.

The results from Coffman et al. (2019) suggest that providing liquidity to prospective teachers could potentially allow individuals to join the teaching profession. But before concluding that easing liquidity constraints will be an effective tool to generate more career teachers—and potentially ease the teacher shortages described above—a few additional questions remain.

First, the prior work looks at the decision to join TFA as measured by whether an individual is teaching as part of the TFA program on the first day of school of the first year. TFA is a two-year program, however, and there is some evidence of attrition out of the program over time (Coffman et al., 2017). One question is whether the marginal teachers induced into the program by the additional liquidity end up dropping out of the program or whether they make it through the two-year commitment.

Second, TFA is only one way to become a teacher and it may be costlier than other routes into teaching available to individuals (e.g., TFA provides TGL funding in part because it asks teachers to travel to get trained during the summer and regularly places them in jobs in new, often expensive cities). Consequently, another question is whether those in the control group who do not join TFA end up becoming teachers through other (perhaps less-costly) channels. Coffman et al. (2019) used survey data to answer this question and found evidence that teachers induced into TFA by the marginal liquidity were mostly pulled out of private sector jobs, rather than out of other teaching jobs.[2] That said, there were some limitations with what could be concluded from that data, as is often the case with survey data. The response rate for the survey was 52.5% for the relevant population. While relatively high compared to many surveys, it may be that the data are missing many non-TFA teachers (e.g., perhaps those who wanted to join TFA but could not because of financial constraints). In addition, due to the timing of the survey, some of the responses

[2]Initial career placement is important both for long-run earnings and industry-placement (Kahn, 2010; Oreopoulos et al., 2012; Altonji et al., 2016; Zhang and de Figueiredo, 2018)

were aspirational, asking, for example, what individuals plan to be doing in two years.

Third, Coffman et al. (2019) could not answer the question of whether teachers—including those induced to join TFA by the marginal liquidity—remain in the teaching profession after their two-year commitment. When considering using liquidity to address the teacher shortage, one might be particularly concerned whether it can induce individuals into teaching in the medium run.

This paper answers these questions by complementing the data from Coffman et al. (2019) with additional data sources. We received additional administrative data from TFA on whether teachers dropped out of TFA before the end of the two-year commitment (and, if so, when). We also conducted a large-scale data collection effort targeting publicly available data on the Coffman et al. (2019) study subjects, including those who did not join TFA. Specifically, we hired a team of seven research assistants to search for public data about each of the 7,295 subjects from the original study and to code their labor market and educational outcomes for five academic years (from 2015–2016 through 2019–2020). Since the TGL experiment was run on prospective teachers who were invited to start teaching through TFA in the falls of 2015, 2016, and 2017, this covered at least three academic years from when they were admitted to TFA, which includes at least one year after the two-year TFA commitment ended. This data collection endeavor took 13 months (from June 2020 through June 2021) and yielded data on 6,036 subjects.

Combining this new administrative and publicly available data with the administrative and survey data from Coffman et al. (2019) allows us to extend the findings from that prior work and answer the questions raised above.

We generate three new findings. First, using the new TFA administrative data, we find that the effects of getting individuals to become teachers with additional liquidity persists through the two-year program, even as the loans provided by the TGL program are required to be paid back.[3] Among the group with the highest financial need, we find that each extra $100 in liquidity increases the probability that an individual starts teaching for TFA by 1.8 percentage points; the same $100 in liquidity increases the probability of completing the two-year commitment by 1.53 percentage points or 85% of the original effect.

Second, we estimate that the extra $100 in liquidity increases the likelihood that an individual is teaching in the first year after being offered the funds—through TFA or otherwise—by 1.17 percentage points. This treatment effect represents 65% of the 1.8 percentage point treatment effect that we measure for starting to teach through TFA. This suggests that some members of the control group did not join TFA because of liquidity constraints still find their way into teaching in that first year, but providing extra liquidity still increases the number of individuals who become teachers.

Third, we find that there is still a (marginally) statistically significant treatment effect of each $100 of liquidity on whether individuals are teaching in the second year after being offered the funds—through TFA or otherwise—of 0.70 percentage points. This treatment effect represents 60% of the first-year effect of becoming a teacher of 1.17 percentage points. This finding suggests

[3] The loans provided by the TGL program are zero-interest loans that teachers are expected to pay back in 18 equal installments starting 6 months after they begin teaching (so that the loan is repaid by the end of their two years of teaching).

that providing liquidity still has an impact in year two, but that more of the control group finds their way into teaching when given more time. The estimated effect is comparable, but no longer significantly significant, at 0.44 percentage points, in the third year after funds are offered. So this point estimate suggests the possibility that liquidity has encouraged individuals to teach even beyond the two-year TFA program, but this result is highly speculative, not just because of the lack of statistical significance but because these estimates rely solely on publicly available data and survey responses rather than than TFA administrative data, which ends after two-years.

Our paper proceeds as follows. Section II describes other policies aimed at increasing teacher recruitment and retention as well as research related to the effects of liquidity on employment. Section III provides background on Teach for America and its Transitional Grants and Loans program. Section IV describes the experiment we leverage in this paper and the empirical strategy we deploy. Section V describes the three sources of data we use for the analysis in this paper. Section VI describes our results. Section VII discusses policy considerations and concludes.

II. RELATED LITERATURE

POLICIES AIMED AT RECRUITING TEACHERS

A number of policies have been implemented with the goal of attracting and retaining more teachers into the profession. The federal Teacher Education Assistance for College and Higher Education (TEACH) grant program provides grants to students who are enrolled in educational programs preparing them to become teachers and who commit to teach for four years (US Dept of Ed, 2023a). Similar programs exist at the state level to provide financial support for students who are training to become teachers. In addition, federal and state programs provide loan forgiveness for individuals who have taught for a certain number of years (US Dept of Ed, 2023b). Some such programs also provide bonuses for teachers who teach in certain schools (see, e.g., Feng and Sass (2018) on the Florida Critical Teacher Shortage Program, Clotfelter et al. (2008) on the North Carolina Bonus Program, and Steele et al. (2010) on the California Governor's Teaching Fellowship program).[4]

Less common are policies that target upfront funding to individuals who are transitioning into teaching, such as by providing funds when individuals may have extra expenses related to taking on the new job (see Liu et al. (2004) on Massachusetts Signing Bonus Program for a notable exception). However, Coffman et al. (2019) suggests that one barrier for a particular subset of teachers is liquidity constraints that prevent them from making the kinds of short-term investments that may be necessary to become teachers. If these types of liquidity constraints are indeed binding, we can improve policies aimed at recruiting teachers in two ways. First, policies could provide prospective teachers with financial help even before they beginning teaching (e.g., upon signing up to teach). Second, policies could provide more of this help at the same cost by providing a larger share of it as loans (i.e., rather than grants), since results show that any form that provides liquidity may be equally effective.

[4]Note while many policies are aimed at recruiting lower-SES students into schools (Andrews et al., 2020) and other policies aimed at recruiting teachers into lower-SES districts, policies aimed at recruiting low-SES teachers remain scarce.

UI AND GRANTS AND LOANS FOR STUDENTS

There are a few lines of research in which liquidity is not the main focus, but where it could potentially contribute to results on labor supply.

Unemployment insurance (UI) provides funds for individuals who may take advantage of the liquidity provided by UI to find a better job, but there is little consensus on the impact of UI on this margin. For example, while both agree UI increases the duration of job search, Herkenhoff et al. (2016) finds that UI increases the pay of the eventual job match while Card et al. (2007) finds no evidence that UI impacts the quality of the eventual job match.[5]

Financial support during college—including through Pell grants and student loans—could impact student outcomes in a number of ways, including by giving students extra liquidity. Pell grants and student loans have been shown to increase graduation rates and future earnings (Black et al., 2020; Denning et al., 2019). In addition, higher student loan debt has been shown to dissuade students from taking jobs in public service, highlighting that debt obligations can also have an impact on individuals' labor market decisions (Field, 2009; Rothstein and Rouse, 2011; Zhang, 2013).

III. BACKGROUND: TFA AND TRANSITIONAL GRANTS AND LOANS

TFA is a nonprofit that places teachers in schools in low-income communities across the United Sates. TFA trains and supports these teachers, but they are otherwise regular school employees. Prospective TFA teachers apply between September and April, attend a six-week training program in the summer, and

[5] For more research on the impact of UI, see Centeno and Novo (2006); Ours and Vodopivec (2008); Addison and Blackburn (2000).

begin teaching in the subsequent academic year. TFA teachers are expected to remain in the program for two years, and TFA estimates that more than half stay beyond the two-year commitment. TFA has a selective admissions process and recruits at hundreds of colleges and universities, including many that are highly ranked.

TFA offers the Transitional Grants and Loans (TGL) program to help cover the costs of transitioning into the new teaching job. To apply for the funding, prospective teachers must complete an extensive application, which includes providing financial information and related documentation. They can apply for the TGL program at any point during the application process. TFA aims to provide TGL offers at the time the applicant is accepted into the program or very soon thereafter (acceptance to TFA is independent of potential TGL need).

In the years of our study, the package of grants and loans offered to an applicant from the TGL program was determined by two key variables. The first is the applicant's "expected expense," which is how much TFA estimates an applicant will need to spend to move to the city they have been placed by TFA and to travel and finance themselves during the TFA summer training. The second is the applicant's "expected contribution" (EC), which is how much TFA estimates an applicant can afford to contribute to the aforementioned expenses.

In the years of our study, the TGL program constructed a package of grants and loans such that the sum of funds offered was equal to the applicant's expected expense minus their expected contribution. However, TFA imposed a rule that TGL packages could not exceed expected expense. As a result, applicants estimated to have a negative expected contribution (e.g., if their outstanding credit card debt exceeded their liquid assets) might not have received enough grants and loans to cover all of the expenses associated with becoming a teacher through TFA. Roughly 10% of TGL applicants—the bottom decile—had negative EC and thus fall in this category.

Almost all TGL funds are disbursed in late May and June before applicants begin their summer training. While the funds are intended for expenses associated with transitioning into teaching, such as travel, moving, testing, and certification fees, the use of the funds is not restricted or monitored. Grants are unconditional and loans are interest free and recipients are asked to repay them starting in January of the first year of teaching (the standard repayment schedule is for 18 equal-sized monthly payments to be made over the subsequent eighteen months). A TGL package typically consists of both a grant and a loan, with lower EC applicants having grants comprise a larger portion of their TGL package.

IV. DESIGN

IV.A. Treatments

Coffman et al. (2019) reports on an experiment run with the TGL program over three years. It involved 7,295 individuals who applied to the program in anticipation of beginning teaching in 2015, 2016, or 2017. In these years, a baseline award package was constructed for each applicant using a modified version of the TGL formula. This baseline award was given to applicants in the control group.[6] Applicants randomized to treatment groups received somewhat

[6] The average baseline award was around $4,000, roughly evenly split between grants and loans. However the 10% of subjects with the highest financial need received average baseline

larger TGL packages.

In the first year, applicants were randomized into a control group or one of two treatment groups, with a one-third probability of each. The control group was offered the baseline TGL package, while the two treatment groups were offered a package with an additional $600 in grants or an additional $600 in loans. Midway through the second year of the experiment, a fourth experimental arm was introduced, where applicants received an additional $1,200 in grants beyond the baseline package. The third year of the experiment was run as a self-replication in which the design was unchanged for applicants in the first two deciles of expected contribution (to replicate treatment effects in the first two years of the study), and modified to stress test the null results for the other eight deciles of EC. For these groups, we lowered baseline awards and our two treatments increased TGL packages by $1,800 in grants or $1,800 in loans.

Across the three years of our study, around 35 million dollars—roughly evenly split between grants and loans—were offered to the TGL applicants in our experiment.

IV.B. Empirical Strategy

While the experiment in Coffman et al. (2019) randomized both grants and loans, it found that, for the highest-need group, the effects of additional funding on joining TFA were just as large whether they were offered additional grant funding or additional loan funding. Consequently, for our main analysis here, we collapse grants and loans and ask how extra liquidity (provided as either

awards of around $5,000, with around $3,200 coming from grants and $1,800 in loans.

grants or loans) affects our outcomes of interest.

In addition, the experiment in Coffman et al. (2019) found that the effects of liquidity on the highest-need group were rather linear (the treatment effect of providing an additional $1,200 in grants was about twice the size of the treatment effects observed from providing an additional $600 of grants or loans). Consequently, we follow one of the specifications in that paper that estimates the effect of providing extra liquidity in hundreds of dollars, combining the variation from all treatments.

Because we are interested in the effect of liquidity on the group that has been shown to be impacted by it (i.e., the highest-need group), we follow Coffman et al. (2019) and report our estimates separately for that group and for everyone else. Consequently, our regression specification is:

$$\begin{aligned} Outcome_i = &\beta_1 \cdot ExtraLiquidity(\$100)_i \cdot HighestNeed_i + \\ &\beta_2 \cdot ExtraLiquidity(\$100)_i \cdot NotHighestNeed_i + \\ &\varphi_1 \cdot HighestNeed_i + \varphi_2 \cdot NotHighestNeed_i + \sum_j \gamma^j \cdot Batch_i^j + \boldsymbol{\delta} \cdot \mathbf{X}_i + \varepsilon_i \end{aligned}$$

Where $Outcome_i$ is an outcome (e.g., joining TFA, being observed teaching in a given year, etc.) for individual i. $ExtraLiquidity(\$100)_i$ reports how many hundreds of dollars of extra liquidity individual i was offered in the experiment. $HighestNeed_i$ is a dummy for individual i being in the highest financial need group and $NotHighestNeed_i$ is a dummy for not being included in that group.

The regression always controls for the batch in which the individual was randomized, j, since that is the point of randomization.[7] It also includes additional demographic controls in $\mathbf{X}_i$. ε_i is an error term.

V. DATA

Data from the experiment comes from three sources. We describe the administrative data from Teach For America in Section V.A. We describe the survey data, collected in 2018, in Section V.B. We describe the data from publicly available sources, collected in 2020 and 2021, in Section V.C.

V.A. TFA data

As described above, Teach For America provided financial information about all the prospective teachers who applied for the TGL program, which is how we identified the financial need of the individuals (e.g., allowing us to classify whether they were in the $HighestNeed_i$ group) and how we constructed their baseline TGL package, to which we added additional liquidity (i.e., in the form of extra grants or loans) if they were randomized to one of our treatment conditions.

In addition, for each individual who we randomized, TFA provided outcome data on their progression through the TFA program. In particular, this included whether they were teaching through TFA on the first day of school of the first year, the first day of the spring semester of the first year, the first day of the second year, the first day of the spring semester of the second year, and whether they completed their full two-year commitment. In Section VI.A we

[7] Applicants were randomized into treatments multiple times each year so there are many batches in each year of the experiment. For additional details, see Coffman et al. (2019).

will explore the impact of liquidity on whether individuals reach each of these milestones.

This administrative data is complete for all individuals in our sample, but it does not give us an indication of what individuals do if they do not join TFA, if they leave TFA partway through the program, or once they have completed their two-year assignment. In particular, we cannot observe if these individuals are teachers outside of TFA or if they are in some other industry. For these reasons, we turn to the survey data and publicly available data described next.

V.B. Survey data

In May 2018, TFA emailed a survey to all the individuals in our sample. The survey, which is discussed in detail in Coffman et al. (2019)—including in Section II.A. of its Online Appendix—had two purposes. One was to ask about access to credit markets to establish the liquidity mechanism that is the focus of that paper. The other was to establish what individuals who chose not to join TFA were doing in various years. In particular, all cohorts were asked what they were doing in the first academic year after they applied for TGL funding (i.e., whether they were teaching or working in some other industry). This first academic year was 2015–2016 for the first cohort, 2016–2017 for the second cohort, and 2017–2018 for the third cohort. We also asked the first cohort what they were doing in the third academic year after they applied for TGL funding (i.e., 2017–2018, the year after their two-year TFA commitment would have ended if they had joined TFA).[8] In this paper, we leverage these responses to

[8]The other two cohorts were asked what they were expecting to do in the third year after they applied for TGL funding, which had not yet occurred. Because these responses reflect prospective guesses rather than outcomes, we treat them differently in the analysis that

the survey to complement our outcome data from TFA.

As with any survey, we worry about non-response and whether it varies by treatment. We made a number of attempts in the survey design to mitigate these concerns. First, the survey was framed as providing data to researchers at Wharton who wanted to learn more about the TGL program to help encourage individuals who chose not to join TFA to respond.[9] Second, we offered incentives for all individuals who completed the survey. Since we were particularly interested in the highest-need group, we offered them dramatically larger incentives (either $20 or $40 amazon gift cards for completion), but we also offered incentives for the rest of the individuals (either a 0.5% or 1% chance of receiving a prepaid debit card with $500 on it).[10]

These tactics generated rather high response rates of 52.5% for those in the highest-need group and 36.8% for others. While response rates were higher for those who joined TFA (40.6%) than those who did not join TFA (32%), they did not differ by whether individuals were in the control group (38.4%) or received extra liquidity (38.5%). While these response rates are rather high for an email survey, we did not receive responses from nearly half of those in the highest-need group, and there are many academic years of work that we do not observe in the survey, particularly for the latter two cohorts. Consequently,

follows. At the end of the survey, we also invited individuals to provide a LinkedIn profile if they had one, which we use to validate some of the publicly collected data as described in Appendix A.I.

[9] TFA believed that teachers were more likely than non-teachers to respond to their surveys and so having an outside research team, particularly one from a business school, might help encourage non-teachers to respond.

[10] The variation in incentives within each group was implemented randomly in the hopes of identifying any selection bias in survey response. Section II.A. of the Online Appendix of Coffman et al. (2019) reports limited differences between those who respond to the higher incentives and lower incentives within each group, suggesting limited selection.

we complement this data with the publicly available data described next.

V.C. Publicly available data

From June 2020 to June 2021, a team of seven research assistants (RAs) conducted internet searches to find education and labor market information on our study subjects. RAs were provided with a set of identifiers including an individual's full name, undergraduate institution, college graduation year, graduate school, and graduate school graduation year (if they attended).[11]

Given this information on each individual, RAs looked for publicly available data on their employment for academic years 2015–2016 through 2019–2020 inclusive. Initial exploration suggested a particularly efficient protocol for identifying such data. First, the RA investigated whether the individual had a LinkedIn profile. Because LinkedIn provides a public social media platform where individuals post their employment and education history, it was rather easy to find a profile that matched the available information (i.e., name, college, and graduation year) if such a profile existed.[12] If the RA could not find an individual on LinkedIn, or they could not find all the desired information on their LinkedIn profile, they followed a search process to find other publicly available data about them, such as on Twitter profiles, public Facebook profiles, teacher directories, and classroom websites, to piece together as complete an employment history as possible (see additional details in Appendix A.I).

[11] These data were unlinked from the other TFA data provided to us (e.g., they were not associated with their financial data or measures of their financial need and they were not associated with outcome data on whether or not the individual joined TFA or took our survey).

[12] Because of the value of LinkedIn as a source of data, we built a custom web scraper to match identifying information with LinkedIn profile urls to provide a natural starting point for each search. For more information on the scraper criteria, as well as additional information on the search process, see Appendix A.I.

After identifying sources of employment information, the RAs coded information about an individual's employment and schooling for each of the five academic years within the window. The primary goal was to identify, in each of the five academic years, whether or not an individual was a teacher. If an individual was a teacher in a given year, the type of school they taught at was recorded. If an individual was not a teacher, the industry of their non-teaching job was recorded. In addition, additional schooling and degrees earned were recorded.[13]

This process yielded at least some employment information on 6,036 of the 7,295 individuals (i.e., 85% of individuals in our sample). Upon the completion of the process, the coded publicly available data—absent the information we used to identify individuals—was merged back in with our other study and administrative data.

VI. RESULTS

In this section, we analyze the impact of providing additional liquidity on our various sources of outcome data. In Section VI.A, we explore whether the additional liquidity has persistent effects on participation as a teacher through the two-year TFA program. In Section VI.B, we combine our various sources of outcome data to investigate the impact of liquidity on whether individuals were teaching for up to three academic years after being offered the additional

[13] RAs were instructed to give confidence ratings for each search they conducted, based on the availability of information for an individual across the five-year period and their perception of the credibility of the sources that they did find. Given the nature of data available on the Internet and on LinkedIn profiles, some individuals only have career information for certain years within the five-year period. Others have complete information. This information is reflected in the individual-level confidence ratings. Figure A-1 shows the distribution of these confidence ratings.

liquidity as part of the TGL experiment. The results in Section VI.B allow us to assess the impact of extra liquidity on generating new teachers who stay in the profession beyond their two-year commitment to TFA.

VI.A. The Effect of Liquidity on Progressing through TFA

Coffman et al. (2019) focused on how marginal liquidity (provided either as grants or loans) impacted the decision to join TFA, as measured by teaching on the program's first day of school (i.e., the first day of the first year of the two-year commitment). Here we extend those findings by exploiting additional TFA administrative data on each TGL applicant teacher, including whether they made it through the full two-year commitment and when they dropped out if they did so.

As described in Section IV.B, to present the most precise estimates of marginal funding, our regressions combine variation from all the experimental treatments that provided additional liquidity to study subjects. In our tables, the coefficient on *Extra Liquidity ($100s)* reports the estimated effect of offering an additional $100 in liquidity to prospective teachers.

The first Panel of Table I demonstrates the effect of additional liquidity on an individual teaching through TFA on the first day of school (i.e., "Fall Y1") through several benchmarks during the two-year program. The table shows that the highest need individuals are 1.80 percentage points more likely to begin teaching with TFA for every $100 in additional liquidity they are offered through the TGL program ($p < 0.001$), a result which is replicated from Coffman et al. (2019).[14] Our new administrative data shows that this

[14]Consistent with results Coffman et al. (2019), and as shown in Appendix Table II, results

effect on teaching as part of TFA persists through the various milestones in the TFA program, including beginning teaching in the second semester of the first year (i.e., "Spring Y1"), teaching on first day of the second year (i.e., "Fall Y2"), teaching in the second semester of the second year (i.e., "Spring Y2"), and completing the two-year program (i.e., "Complete"). The effect persists at nearly its full size over time, and there is still a 1.53 percentage point impact on completing the two-year program for every $100 in additional liquidity a prospective teacher was offered. This represents 85% of the 1.80 percentage point effect of beginning to teach, suggesting the effect persists many years after the additional liquidity was offered. As in Coffman et al. (2019), we observe no effect of liquidity on those who are not in the highest need group.

Understanding the lasting impacts of upfront grants and upfront loans separately is also important. In showing positive and similarly sized effects for upfront grants and loans, Coffman et al. (2019) are able to conclude that it is liquidity that matters for getting teachers in the door on day one. Are the lasting effects in Table III due to liquidity? It might be that grants and loans are differently effective over time. For example, do some cash-constrained teachers drop out when they have to repay their loans?[15]

Table II shows the separate, lasting effects of grants and loans on progressing through the first two years of teaching through TFA. The table shows that it is upfront liquidity that matters for the full two years. Upfront grant money provided to the highest need group produced substantially more teachers throughout the two-year commitment, through completion, with treatment

are similar when considering the two forms of liquidity (i.e., grants and loans) separately.

[15]Recall loan-repayment starts in January of the second semester of the first year

TABLE I
EFFECTS ON TFA PARTICIPATION

	Replication	New Data			
	Fall Y1	Spring Y1	Fall Y2	Spring Y2	Complete
Extra Liquidity ($100s)	1.80***	1.70***	1.58***	1.54***	1.53***
× Highest Need	(0.40)	(0.42)	(0.45)	(0.45)	(0.46)
Extra Liquidity ($100s)	0.01	0.02	-0.04	-0.01	-0.01
× Not Highest Need	(0.10)	(0.10)	(0.11)	(0.11)	(0.11)
Highest Need	-11.82***	-11.41***	-12.45***	-13.01***	-12.25***
	(2.65)	(2.79)	(2.95)	(2.98)	(3.02)

Notes. Table shows Linear Probability Model (OLS) regressions of whether an applicant joins TFA from the regression specification described in Section IV.B. Robust standard errors are reported in parentheses. *, **, *** denote $p < 0.10$, 0.05, and 0.01, respectively. All regressions include a set of demographics that include: a linear age term; a linear term for the applicant's "fit" with TFA, as measured as part of the application process; and dummies for race, gender, assigned region, whether the applicant was assigned to her most preferred region, and whether the applicant was assigned to her most preferred subject. We also include a missing data dummy for each demographic variable that is sometimes missing (age, race, and fit). Additionally, all regressions include fixed effects for the batch in which the applicant's TGL awards were processed, the point at which randomization occurred.

effects ranging from 1.49–1.64 percentage points per $100 of upfront grant money provided to the most financially constrained ($p < 0.01$ throughout). The estimates are similar, albeit directionally larger, for upfront loan money provided to the highest need group. Treatment effects range from 1.62–2.24 percentage points per $100 in upfront loan offers ($p < 0.01$ through Spring Year 2 and $p < 0.05$ for completing their TFA commitment). Either policy is effective at allowing financially constrained candidates to become teachers, and the estimates suggest they are equally effective. Liquidity produces TFA teachers, even two years later.

However, as noted above, these results only account for individuals becoming teachers through TFA and only follow individuals for two years. Individuals may become teachers through other channels, and we are also interested in whether individuals remain in the teaching profession after the TFA program ends. Consequently, in the next section, we report results including our additional outcome data.

VI.B. The Effect of Liquidity on Teaching Anywhere

To analyze the impact of additional liquidity on the likelihood of teaching anywhere—both within TFA or outside of TFA—we construct a measure of teaching that combines our data sources. We start with the assumption that no one is teaching in any academic year unless we find affirmative evidence that they are teaching in that academic year. We start with the TFA data, and we say that an individual is teaching in an academic year if TFA reports that they taught at all during that academic year. For anyone who is not in the TFA data for a given academic year, we say that they are teaching if they tell us

TABLE II

EFFECTS ON TFA PARTICIPATION FOR GRANTS AND LOANS SEPARATELY

	Replication	New Data			
	Fall Y1	Spring Y1	Fall Y2	Spring Y2	Complete
Extra Grants ($100s) × Highest Need	1.77***	1.64***	1.53***	1.49***	1.52***
	(0.41)	(0.43)	(0.45)	(0.46)	(0.46)
Extra Grants ($100s) × Not Highest Need	0.05	0.04	-0.00	0.04	0.04
	(0.11)	(0.12)	(0.12)	(0.13)	(0.13)
Extra Loans ($100s) × Highest Need	2.06***	2.24***	2.04***	1.97***	1.62**
	(0.65)	(0.69)	(0.73)	(0.73)	(0.74)
Extra Loans ($100s) × Not Highest Need	-0.05	-0.02	-0.10	-0.09	-0.08
	(0.12)	(0.13)	(0.13)	(0.13)	(0.14)
Highest Need	-12.15***	-12.07***	-13.01***	-13.56***	-12.39***
	(2.72)	(2.87)	(3.04)	(3.06)	(3.10)

Notes. Table shows Linear Probability Model (OLS) regressions of whether an applicant joins TFA from the regression specification described in Section IV.B. Robust standard errors are reported in parentheses. *, **, *** denote $p < 0.10$, 0.05, and 0.01, respectively. All regressions include a set of demographics that include: a linear age term; a linear term for the applicant's "fit" with TFA, as measured as part of the application process; and dummies for race, gender, assigned region, whether the applicant was assigned to her most preferred region, and whether the applicant was assigned to her most preferred subject. We also include a missing data dummy for each demographic variable that is sometimes missing (age, race, and fit). Additionally, all regressions include fixed effects for the batch in which the applicant's TGL awards were processed, the point at which randomization occurred.

in our survey that they were teaching in a given academic year. We also code them as teaching in an academic year if we see evidence of them teaching in the publicly available data (e.g., they list a teaching position on LinkedIn or we see them listed on a school's website as a teacher for a given academic year).[16] Finally, we recode them as *not* teaching if the survey data are aspirational (i.e., "I plan to be a teacher" as opposed to "I am currently teaching"), they are found in the publicly available data for that year, and the public data say they are not teaching.[17] In short, we code them as teaching if any data source says they are teaching that year, unless that data source is aspirational survey data and we have other data that refute it. We also consider a paradigm that codes everyone as a teacher if any data source says they are teaching, regardless if the data are aspirational or determined (Table A-2 in the Appendix).

There are two reasons why individuals might not be observed to be teaching in a given academic year. The first is that they were observed in our data but did not teach in that year (e.g., they were employed elsewhere, were in a full-time educational program, etc.). The second is that they were not in our data because we failed to find them in that academic year (i.e., they were not in TFA, did not respond to our survey or were not asked about a particular academic year in our survey, and we could not find their occupation, teaching or otherwise, in the publicly available data for that academic year). Consequently,

[16] Note that we label individuals as teaching in a given academic academic year if there is any indication that they are teaching in that academic year (i.e., we do not attempt to parse whether individuals teach for only part of the year). Results from Table I suggests limited scope for individuals teaching for partial years.

[17] Overriding the aspirational survey data only affects the estimates for the third year and only for the latter two cohorts of the study. If we do not recode these individuals and assume they were teachers, we would get an estimate of the effect of liquidity that is directionally larger for the highest need group (i.e., 0.59) but also not quite statistically significant.

confidence in our treatment estimates—about whether additional liquidity induces individuals to teach in a given year—depends on whether the treatment makes individuals more likely to appear in our data at all.[18]

Consequently, in Table III we report the impact of treatment on both the likelihood of being a teacher in a given academic year based on the definition above and on whether we find them in our data (whether they are "Found"). If extra liquidity makes someone more likely to be found in our data, we might worry that our estimated treatment effect could be biased upward, since to be classified as teaching in a given year, an individual needs to be found. Similarly if extra liquidity makes someone less likely to be found in our data, our treatment effect could be biased downward.

Table III shows the results from our main specification for each of the three academic years after an individual is offered funding through TGL. The first column for Year 1 shows that for each extra $100 in liquidity offered to individuals by the TGL program, individuals are 1.17 percentage point more likely to be observed teaching in the first academic year after being offered their TGL package. This 1.17 percentage point effect represents 65% of the estimated effect on beginning teaching for Teach For America (the 1.80 percentage points shown in the first column of Table I). This suggests that while some of the highest need individuals who are not offered extra liquidity find their way into teaching through other channels, the liquidity indeed generates additional teachers in that first year.

[18]This is analogous to the argument for why we cannot simply rely on TFA data to establish an increase in teaching overall; our TFA data only observes people teaching through TFA, so when we get a treatment effect on being a teacher through TFA this is the same as a treatment effect on appearing in our data at all.

TABLE III

EFFECTS ON TEACHING ANYWHERE (AND ON EMPLOYMENT DATA BEING FOUND)

	Year 1		Year 2		Year 3	
	Teaching	Found	Teaching	Found	Teaching	Found
Extra Liquidity (\$100s) × Highest Need	1.17^{***}	0.15	0.70^{*}	0.11	0.44	-0.12
	(0.36)	(0.27)	(0.40)	(0.34)	(0.48)	(0.39)
Extra Liquidity (\$100s) × Not Highest Need	0.02	-0.01	0.00	-0.05	-0.11	-0.08
	(0.09)	(0.07)	(0.10)	(0.08)	(0.12)	(0.10)
Highest Need	-7.19^{***}	2.48	-5.14^{*}	-0.87	2.14	7.80^{***}
	(2.40)	(1.82)	(2.66)	(2.27)	(3.15)	(2.59)
TFA Admin Data	✓	✓	✓	✓		
Survey Data	✓	✓			✓	✓
Public Data	✓	✓	✓	✓	✓	✓

Notes. Table shows Linear Probability Model (OLS) regressions of whether an applicant is observed teaching in each academic year ("Teaching") and whether they are found at all (i.e., teaching or otherwise) in each academic year ("Found") from the regression specification described in Section IV.B. Robust standard errors are reported in parentheses. *, **, *** denote $p < 0.10$, 0.05, and 0.01, respectively. from the regression specification described in Section IV.B. Year 3 estimates do not include TFA data, as it is not available. All regressions include a set of demographics that include: a linear age term; a linear term for the applicant's "fit" with TFA, as measured as part of the application process; and dummies for race, gender, assigned region, whether the applicant was assigned to her most preferred region, and whether the applicant was assigned to her most preferred subject. We also include a missing data dummy for each demographic variable that is sometimes missing (age, race, and fit). Additionally, all regressions include fixed effects for the batch in which the applicant's TGL awards were processed, the point at which randomization occurred.

The second column for Year 1 shows that individuals are not statistically significantly more likely to be found when they are offered more liquidity. And while the point estimate is positive (0.15), it is small (only 13% of the estimated coefficient on Teaching in that year). Even if we assumed that all of the extra people who were found in the treatment group were teaching—and we wanted to remove them from our estimates to be conservative—we would still estimate that each extra $100 of liquidity increases the probability of teaching of 1.02 percentage points (i.e., 1.17–0.15). This effect would remain significant and suggests that $10,000 of liquidity (e.g., in the form of two-year, interest-free loans as offered by TFA) induces an extra person to become a teacher in expectation.

It is also worth noting that the coefficient on the highest need group in the Teaching column of Year 1 is negative, significant and large, suggesting that highest need individuals who are not offered extra liquidity are dramatically less likely to become teachers than those who are better off financially, even conditional on them all applying to and being offered jobs by TFA.[19]

The first column for Year 2 shows that liquidity still has an impact two years later. The coefficient estimate is somewhat smaller, at 0.70 percentage points per $100 of liquidity. This attenuation of the treatment effect into the second year (i.e., 0.70 is 60% of the 1.17 effect observed in Year 1) suggests that some of the individuals not offered extra liquidity manage to make their way into teaching by the second year. This is consistent with the coefficient

[19]The highest need individuals in the control group are also somewhat more likely to be found than their counterparts with less financial need, possibly because of the higher incentives offered to the highest need individuals to take our 2018 survey.

on the highest need group also decreasing, suggesting that the highest need individuals who are not offered extra liquidity catch up somewhat with their lower need counterparts by the second year.

Finally, the second-to-last column in Table III shows that the impact of liquidity is estimated to be positive three years later, although the 0.44 treatment effect is no longer statistically significant ($p = 0.35$). The -0.12 estimate on being found in Year 3 suggests the 0.44 may be biased downward (as this group is less likely to be found, and those not found in our data are assumed to be not teaching). While insignificant, the similarity of coefficients across Year 2 and Year 3 gives the impression that some of the increase in teaching due to the extra liquidity may continue into the medium term.

When interpreting the results from Year 3 (i.e., the last two columns of the table), it is worth emphasizing that—unlike the results for the first two years, which are buttressed by high-quality administrative data from TFA—the estimates rely *only* on survey data and publicly available data. Though these data sources are very useful supplements to the TFA data, results using these data alone should be interpreted with less confidence and due qualification.[20] Nonetheless, the data do provide a reasonable, even if noisy, snapshot of reality beyond what we can infer from the administrative data.

The Appendix provides some robustness checks of the results. Table A-1 utilizes the estimation strategy as in Table III but does not use public data that the RAs coded as "low confidence" (see details in Appendix A.I). Figure

[20]For example, using these data alone, we would not find a treatment effect of providing liquidity to the highest need individuals on joining TFA in the first year. The TFA data are near-perfect measures of who is in TFA, so this is (only) concerning for the supplemental data sources.

A-1 shows the histogram of confidence ratings for all coded public data. The RAs are confident or very confident (i.e., rated as a $2-3$) in a majority of the data, but there is a meaningful portion of low confidence data as well. Whereas Table III uses all public data, Table A-1 only uses data coded with a confidence of 2 or 3, excluding data whose confidence was coded as 0 or 1. Effectively, this uses less publicly available data, and thus recodes many from being teachers to being non-teachers (as being a non-teacher is the default in our coding paradigm). Using these data only, the effect of providing liquidity to the highest need group is substantial and significant both in the first year (1.22 percentage points per \$100) and in the second year (0.72 percentage points per \$100). The effect on the third year is positive (0.41 percentage points per \$100) but insignificant ($p = 0.39$). That the Year 3 estimate is directionally smaller in this specification is to be expected since this year's estimate relies heavily on publicly available data, and this exercise mechanically replaces instances of teaching with not teaching in the data. Consequently, the average treatment effect will shift towards zero.

Table A-2 uses data that recodes anyone as a teacher if any data source indicates they are teaching that year—including aspirational data—and runs the same analysis as Table III. This recoding only affects the Year 3 estimates. The effect of providing liquidity to the highest need group is a 0.60 percentage point increase in the likelihood of teaching three years later for each \$100 of liquidity. This estimate is directionally larger than the Year 3 estimate in Table III, and roughly in line with the Year 2 estimate, but it is still insignificant ($p = 0.21$).

The totality of evidence suggests that offering teachers liquidity in the months before they would begin teaching can increase the number of teachers. We see this effect in the short term—many individuals appear to begin teaching a year or two earlier when they have access to liquidity. This effect may persist somewhat into the medium term (e.g., if those who fail to become teachers due to the lack of financial liquidity never make they way into teaching), although our estimates on that margin are more speculative.

VI.C. Value of publicly available data

In addition to presenting the results in the prior sections, a contribution of this paper is demonstrating the feasibility of using publicly available data to supplement other forms of data to conduct our analyses. As a demonstration of the value of collecting this data, Appendix Table A-3 replicates the structure of Table III but imagines we did not have access to the publicly available data.

Comparing the results in Appendix Table A-3 to those in Table III, we see that without the publicly available data, the coefficients on Found are much larger in magnitude and sometimes significant. In addition, the coefficient estimates on extra liquidity, our main variable of interest, also differ from the estimates in Table III in a way that is consistent with the additional imbalance in who is identified in our data—due to the absence of publicly available data—generating bias in our treatment effects on liquidity.

Taken together, these results suggest the value of adding this additional information in our context. While the data collection effort was time intensive, other researchers may find it valuable in their contexts as well. In that case, the details presented in Appendix A.I about how we collected that data may prove useful.

VII. CONCLUSION

In this paper, we test if recent results suggesting additional liquidity could induce individuals to become teachers (Coffman et al., 2019) persist over time. Using new administrative data from TFA, we show that providing modest increases in liquidity to those in financial need increases the number of teachers quite substantially for the entire two years of the TFA program. Using the TFA data, survey data, as well as newly collected publicly available data (mostly from LinkedIn), we show that the intervention indeed produced new teachers—rather than just shifting teachers into TFA—and this result persists for at least two years, and perhaps beyond.

The importance of this research lies in its ability to demonstrate the powerful impact that targeted financial assistance can have in addressing barriers to entry for individuals who aspire to become teachers. Providing liquidity to those in need does not just increase the number of individuals who are able to start teaching through programs like TFA, but also the number of individuals who are able to pursue teaching as a long-term career. Our results suggest the funding also allows individuals who will eventually become teachers to do so earlier.

A final contribution of our paper is helping to pioneer a method for learning about career paths by leveraging publicly available information on individuals from sources like LinkedIn. The Appendix describes in detail the process we took to gather this data. While time intensive, it was feasible to complete within a year and scaled rather linearly with RA hours. While we were unable to find technology that sped up data collection (at least beyond automating the search process on LinkedIn through our customized scraper), additional advancement in technology might lower the costs to researchers even further.

References

Addison, John T. and McKinley L. Blackburn, "The Effects of Unemployment Insurance on Post-Unemployment Earnings," *Labour Economics*, 2000, *7* (1), 21–53.

Altonji, Joseph G., Lisa B. Kahn, and Jamin D. Speer, "Cashier or Consultant? Entry Labor Market Conditions, Field of Study, and Career Success," *Journal of Labor Economics*, 2016, *34* (S1), S361–S401.

Andrews, Rodney J, Scott A Imberman, and Michael F Lovenheim, "Recruiting and supporting low-income, high-achieving students at flagship universities," *Economics of Education Review*, 2020, *74*, 101923.

Black, Sandra E, Jeffrey T Denning, Lisa J Dettling, Sarena Goodman, and Lesley J Turner, "Taking it to the limit: Effects of increased student loan availability on attainment, earnings, and financial well-being," 2020.

Card, David, Raj Chetty, and Andrea Weber, "Cash-On-Hand and Competing Models of Intertemporal Behavior: New Evidence from the Labor Market," *Quarterly Journal of Economics*, 2007, *122* (4), 1511–1560.

Centeno, Mario and Alvaro A. Novo, "The Impact of Unemployment Insurance Generosity on Match Quality Distribution," *Economic Letters*, 2006, *93* (2), 235–241.

Clotfelter, Charles, Elizabeth Glennie, Helen Ladd, and Jacob Vigdor, "Would Higher Salaries Keep Teachers in High-Poverty Schools? Evidence from a Policy Intervention in North Carolina," *Journal of Public Economics*, 2008, *92* (5-6), 1352–1370.

Coffman, Lucas C., Clayton R. Featherstone, and Judd B. Kessler, "Can Social Information Affect What Job You Choose and Keep?," *American Economic Journal: Applied Economics*, 2017, *9* (1), 96–117.

Coffman, Lucas C, John J Conlon, Clayton R Featherstone, and Judd B Kessler, "Liquidity Affects Job Choice: Evidence from Teach for America," *The Quarterly Journal of Economics*, 2019, *134* (4), 2203–2236.

Dee, Thomas S., "Teachers, Race, and Student Achievement in a Randomized Experiment," *Review of Economics and Statistics*, 2004, *86* (1), 195–210.

Denning, Jeffrey T, Benjamin M Marx, and Lesley J Turner, "ProPelled: The effects of grants on graduation, earnings, and welfare," *American Economic Journal: Applied Economics*, 2019, *11* (3), 193–224.

Feng, Li and Tim R. Sass, "The Impact of Incentives to Recruit and Retain Teachers in 'Hard-to-Staff' Subjects," *Journal of Policy Analysis and Management*, 2018, *37* (1), 112–135.

Field, Erica, "Educational Debt Burden and Career Choice: Evidence from a Financial Aid Experiment at NYU Law School," *American Economic Journal: Applied Economics*, 2009, *1* (1), 1–21.

Herkenhoff, Kyle, Gordon Phillips, and Ethan Cohen-Cole, "How Credit Constraints Impact Job Finding Rates, Sorting and Aggregate Output," *National Bureau of Economic Research Working Paper*, 2016, (w22274).

Kahn, Lisa B., "The Long-Term Labor Market Consequences of Graduating from College in a Bad Economy," *Labour Economics*, 2010, *17* (2), 303–316.

Learning Policy Institute, "Preparing and Supporting Effective Teachers for Low-Income Students," 2018.

_ , "The State of Racial Diversity in the Educator Workforce," 2018.

_ , "The U.S. Teacher Shortage: A Continuing Crisis," 2019.

Liu, Edward, Susan Moore Johnson, and Heather G. Peske, "New Teachers and the Massachusetts Signing Bonus: The Limits of Inducements," *Educational Evaluation and Policy Analysis*, 2004, *26* (3), 217–236.

Oreopoulos, Philip, Till von Wachter, and Andrew Heisz, "The Short and Long-Term Career Effects of Graduating in a Recession," *American Economic Journal: Applied Economics*, 2012, *4* (1), 1–29.

Ours, Jan C. van and Milan Vodopivec, "Does Reducing Unemployment Insurance Generosity Reduce Job Match Quality?," *Journal of Public Economics*, 2008, *92* (3), 235–241.

Rothstein, Jesse and Cecilia Elena Rouse, "Constrained After College: Student Loans and Early-Career Occupational Choices," *Journal of Public Economics*, 2011, *95* (1), 149–163.

Steele, Jennifer L., Richard J. Murnane, and John B. Willett, "Do Financial Incentives Help Low-Performing Schools Attract and Keep Academically Talented Teachers? Evidence from California," *Journal of Policy Analysis and Management*, 2010, *29* (3), 451–478.

U.S. Department of Education, "Teacher Shortage Area Nationwide Listing 2018-2019," 2018.

_ , "Receive a TEACH Grant to Pay for College." https://studentaid.gov/understand-aid/types/grants/teach [Accessed: February 14, 2023].

_ , "Teacher Loan Forgiveness." "https://studentaid.gov/manage-loans/forgiveness-cancellation/teacher [Accessed: February 14, 2023]".

Zhang, Congshan and John M. de Figueiredo, "Are Recessions Good for Government Hires? The Effect of Unemployment on Public Sector Human Capital," *Economics Letters*, 2018, *170*, 1–5.

Zhang, Lei, "Effects of College Educational Debt on Graduate School Attendance and Early Career and Lifestyle Choices," *Education Economics*, 2013, *21* (2), 154–175.

Appendix

TABLE A-1
EFFECTS ON TEACHING ANYWHERE
(AND ON EMPLOYMENT DATA BEING FOUND)
USING ONLY HIGHER-CONFIDENCE PUBLIC DATA

	Year 1		Year 2		Year 3	
	Teaching	Found	Teaching	Found	Teaching	Found
Extra Liquidity ($100s) × Highest Need	1.22***	0.07	0.72*	0.06	0.41	-0.25
	(0.36)	(0.28)	(0.40)	(0.35)	(0.48)	(0.40)
Extra Liquidity ($100s) × Not Highest Need	0.02	-0.01	0.00	-0.02	-0.07	-0.04
	(0.09)	(0.07)	(0.10)	(0.09)	(0.12)	(0.10)
Highest Need	-7.45***	2.90	-5.50**	-0.76	2.70	9.18***
	(2.40)	(1.84)	(2.66)	(2.30)	(3.15)	(2.66)
TFA Admin Data	✓	✓	✓	✓		
Survey Data	✓	✓			✓	✓
Public Data	✓	✓	✓	✓	✓	✓

Notes. Table shows Linear Probability Model (OLS) regressions of whether an applicant is observed teaching in each academic year ("Teaching") and whether they are found at all (i.e., teaching or otherwise) in each academic year ("Found") from the regression specification described in Section IV.B. Robust standard errors are reported in parentheses. *, **, *** denote $p < 0.10$, 0.05, and 0.01, respectively. In contrast to Table III, which include all public data, here we include only data that the RA coders rated as high confidence (a rating of 2 or 3 on a 0-3 scale). Year 3 estimates do not include TFA data, as it is not available. All regressions include a set of demographics that include: a linear age term; a linear term for the applicant's "fit" with TFA, as measured as part of the application process; and dummies for race, gender, assigned region, whether the applicant was assigned to her most preferred region, and whether the applicant was assigned to her most preferred subject. We also include a missing data dummy for each demographic variable that is sometimes missing (age, race, and fit). Additionally, all regressions include fixed effects for the batch in which the applicant's TGL awards were processed, the point at which randomization occurred.

TABLE A-2
EFFECTS ON TEACHING ANYWHERE
PRIVILEGING ASPIRATIONAL SURVEY DATA

	Year 3 Teaching
Extra Liquidity ($100s) × Highest Need	0.60
	(0.47)
Extra Liquidity ($100s) × Not Highest Need	-0.15
	(0.12)
Highest Need	2.48
	(3.14)
TFA Admin Data	✓
Survey Data	✓
Public Data	✓

Notes. Table shows Linear Probability Model (OLS) regressions of whether an applicant is observed teaching in each academic year from the regression specification described in Section IV.B. Robust standard errors are reported in parentheses. *, **, *** denote $p < 0.10$, 0.05, and 0.01, respectively. from the regression specification described in Section IV.B. In contrast to Table III, here we count a subject as teaching if they report expecting to teach two years after their initial TFA commitment even when they do not appear to be teaching in the public data. All regressions include a set of demographics that include: a linear age term; a linear term for the applicant's "fit" with TFA, as measured as part of the application process; and dummies for race, gender, assigned region, whether the applicant was assigned to her most preferred region, and whether the applicant was assigned to her most preferred subject. We also include a missing data dummy for each demographic variable that is sometimes missing (age, race, and fit). Additionally, all regressions include fixed effects for the batch in which the applicant's TGL awards were processed, the point at which randomization occurred.

TABLE A-3
EFFECTS ON TEACHING ANYWHERE
(AND ON EMPLOYMENT DATA BEING FOUND)
EXCLUDING PUBLIC DATA

	Year 1		Year 2	Year 3	
	Teaching	Found	Teaching	Teaching	Found
Extra Liquidity ($100s) × Highest Need	1.63***	0.83**	1.58***	0.57	-0.44
	(0.39)	(0.36)	(0.45)	(0.41)	(0.46)
Extra Liquidity ($100s) × Not Highest Need	-0.06	-0.05	-0.04	0.03	0.08
	(0.10)	(0.09)	(0.11)	(0.10)	(0.11)
Highest Need	-9.99***	-0.82	-12.45***	7.38***	18.37***
	(2.58)	(2.36)	(2.95)	(2.68)	(3.06)
TFA Admin Data	✓	✓	✓		
Survey Data	✓	✓		✓	✓

Notes. Table shows Linear Probability Model (OLS) regressions of whether an applicant is observed teaching in each academic year ("Teaching") and whether they are found at all (i.e., teaching or otherwise) in each academic year ("Found") from the regression specification described in Section IV.B. Robust standard errors are reported in parentheses. *, **, *** denote $p < 0.10$, 0.05, and 0.01, respectively. from the regression specification described in Section IV.B. Year 3 estimates do not include TFA data, as it is not available. All regressions include a set of demographics that include: a linear age term; a linear term for the applicant's "fit" with TFA, as measured as part of the application process; and dummies for race, gender, assigned region, whether the applicant was assigned to her most preferred region, and whether the applicant was assigned to her most preferred subject. We also include a missing data dummy for each demographic variable that is sometimes missing (age, race, and fit). Additionally, all regressions include fixed effects for the batch in which the applicant's TGL awards were processed, the point at which randomization occurred.

FIGURE A-1
Histogram of RAs' confidence ratings
of each coded entry in public data

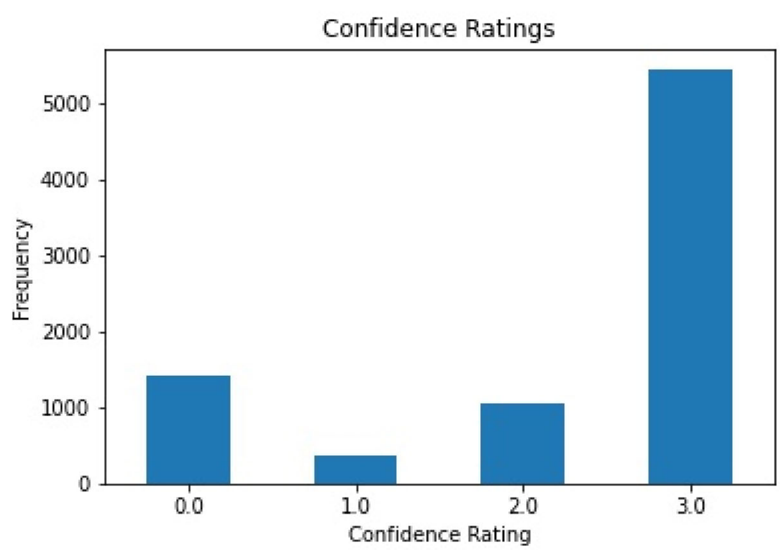

ONLINE APPENDIX

A.I. Appendix: Search Process

A.I.A. Search Process

For each of the individuals in the original experiment, TFA provided us with identifying information to allow RAs to search for the existence of social media profiles or other indications of employment online. For data security and for coding integrity reasons, the information was unlinked from TFA identifiers or any other information we had about individuals (e.g., financial information or treatment status), which guaranteed our coding would be blind to treatment.

Given the set of identifiers about each individual (which included their full name, preferred name, college, college graduation year, graduate school, and graduate school year), research assistants (RAs) were given a set of guidelines about how to locate and record employment information over the five-year period.

To provide a starting point for the RAs, a LinkedIn web scraper first matched the personal information of each individual in the search set with existing LinkedIn profiles, recording a list of URLs that matched the participant search criteria. This was implemented in Selenium, which is a Web Driver Automation Tool, as well as ChromeDriver, which was used to extract information from websites on the Chrome browser. During the initial search process, we noticed that Google's use of Graph Theory in its search engine allowed us to find the existence of a LinkedIn profile hyperlink effectively and efficiently via this scraper because the top two results tended to be the correct person. Using Google was helpful in many ways, but possibly the most important was the ability to maneuver LinkedIn's anti-scraping measures, which risked banning IP addresses. The scraper ran through the identifying information as follows: first, if their legal and preferred names were the same, we disregarded their preferred name. If they were different we included both in the search. We did the same filtering process to college institutions, and then formulated the criteria into the following search format: "legal first name, preferred first name (if applicable), legal last name, institution 1, institution 2 (if applicable), LinkedIn". This text was then put in the Google search bar.

The scraped URLs were provided to the RAs doing the information searching and coding, as a first resource to check, given that LinkedIn profiles tended to have the most robust and reliable employment information over the time period. RAs were instructed to first check the scraped LinkedIn webpage to see if the identifying information was a match. Because these profiles typically display name, college, and date of graduation, it was often immediately obvious whether or not it was the right person. If it was the right person, coding of data would begin. If it was not the right person, the RAs would search LinkedIn for the correct profile. If they did not find one, or if the LinkedIn profile was incomplete (e.g., did not appear up-to-date) they would then perform a more-general Google search. For profiles that had not been updated during the time period, or that showed an individual in the same job for over 5 years, assistants were instructed to find secondary sources to confirm job status during questionable years.

A set of general search keywords such as the individual's first and last name and the word "teacher" or "Teach for America" were used for general Google searches. Such searches, if successful, typically yielded staff directories, wedding pages, Twitter profiles, Facebook profiles, or Open Payroll records, all of which tended to have a date or plausible date range for relevant employment history context. Uncertainty regarding these details were factored into the confidence ratings explained below. The URLs of web pages that matched the individual's identifying information were recorded along with any retrievable information from those pages. If a particular keyword search provided no promising results, assistants were instructed to try a variety of keyword combinations on Google. Assistants were instructed to stop searching and move on to the next participant if they found themselves searching for longer than 10 minutes on one individual with no results.

One RA was assigned to find and code each individual. A total of 7 RAs worked on the process, and found at least some information for 6,036 of the 7,295 individuals in the search set. The process took a total of one year (from June 2020 to June 2021).

A.I.B. Search Accuracy Results

As mentioned in the paper, the May 2018 survey asked respondents to provide their LinkedIn profile url, if they had one. We were able to verify search accuracy for these respondents. There were 575 survey respondents who responded to the question which asked to provide a LinkedIn profile. We directly compared these profiles to the ones which RAs recorded throughout the search process, using string matching on the urls. Our LinkedIn scraper exactly matched 399 of these profiles, while our RAs successfully recovered an additional 130, of which the scraper did not provide a match for or provided an incorrect profile. This suggests a 92% accuracy rate (i.e., 529/575), when restricting only to LinkedIn data. Note that when RAs encountered profiles that were sparse or unavailable, they would supplement the data with external webpages from Google, classroom websites, and other sources, suggesting the rate of finding individuals with publicly available data might be higher accounting for all possible data sources.

A closer inspection of the 46 LinkedIn profiles provided in the 2018 survey that the RAs did not find was conducted in order to understand why RAs missed the profile or recorded a different one. There were 7 profiles that no longer existed (e.g., the person deleted their LinkedIn account). Of the 39 remaining profiles, 10 seemed to have additional security measures that prevented access.[1a] The remaining 29 (5% of the profiles) were failed to be correctly found by the RAs.

A.I.C. Coding Process

RAs looked to find the most comprehensive information possible over the five-year period, with the primary focus being whether or not the individual was a teacher in a given year. In the majority of cases with a LinkedIn match, there was comprehensive information available regarding the individual's employment history over the five-year period, including whether or not they acquired additional degrees, whether they were a teacher, and the type of

[1a]This was apparent when you navigate to the url and LinkedIn reports that it does not register the profile or that it cannot verify that url, as opposed to the profiles that have been deleted, which have a "this profile no longer exists" error message.

school where they worked. Information was coded as follows.

First, RAs coded the individual's primary occupation for each year in the period. These options were:

1. teacher
2. education-related role (e.g., non-teacher roles such as curriculum developer or special needs coordinator)
3. non-education occupation
4. full-time student
5. unemployed

If the person was a teacher, the type of school was recorded, if available, including:

1. public
2. private
3. charter

If in a given year, the individual was in a non-education profession, this was recorded as their primary role, as well as classified into one of the following broad categories:

1. secretarial or back-office roles
2. research or academia roles
3. business-related roles
4. law-related roles
5. other

RAs coded the years in which an individual listed that they were in a graduate degree program along with the type of graduate degree. This data was collected whether or not the person was a full-time student; that is, if a person was simultaneously a teacher and getting an education degree, we coded them as primarily a teacher, and then recorded, from the following options, the type of degree they pursued:

1. teaching and education
2. medical
3. business
4. law
5. other

Finally, the RA provided a confidence rating to reflect the level of accuracy and completion of the individual's records. This confidence was rated on a scale from 3 to 0, with 3 being the highest confidence and completeness of the search process accuracy and recorded search detail, and 0 indicating that the RA was unable to find any information corresponding to the individual.

The coding process for cases with only partial information and less identifiable sources was the same. This uncertainty would be captured by a rating of 1 or 2 in the confidence category. For individuals without a LinkedIn profile, with an incomplete LinkedIn profile, or a lack of comprehensive additional web page sources, the RA coded the available information (e.g., for the subset of the years where information was available). For instance, if an individual had no LinkedIn profile or social media, but was listed as a teacher for two years during the period on an "Open Wages" website for a state's public school teachers, the primary role of teacher would be recorded for only these 2 years. These searches might be quantified as a 1 or 2 on the confidence rating depending on information available throughout the rest of the period.