Liquidity For Teachers: Evidence from Teach For America and LinkedIn*

Lucas C. Coffman
John J. Conlon
Clayton R. Featherstone
Judd B. Kessler
Jessica Mixon

August 14, 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. Nguyen et al. (2022) report conservative estimates of 36,000 teaching vacancies as well as 163,000 positions currently filled by underqualified teachers. 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 (U.S. Department of Education, 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. Moreover, the demographics of classrooms are imbalanced: African American, Hispanic, and American Indian students are more likely to have teachers of a different race than are White students (U.S. Department of Education, 2016). As shown in Dee (2004), having a race-matched teacher improves student performance.[1]

Recent experimental evidence in Coffman et al. (2019) identified a promising 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 work showed

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

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. 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 who did not join TFA because of liquidity constraints still find their way into teaching in that first year, but that 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 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.

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 (U.S. Department of Education, 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 (U.S. Department of Education, 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).

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 begin 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), as our results show that any form that provides liquidity may be equally effective.

III. BACKGROUND: TFA AND TRANSITIONAL GRANTS AND LOANS

TFA is a nonprofit that places teachers in schools in low-income communities across the United States. 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 begin teaching in the subsequent academic year. TFA teachers are expected to remain in the program for two years.

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 that includes financial information and related documentation. They can apply for the TGL program at any point during the application process, though the vast majority do so only after having been accepted into—and formally agreeing to join—the program. Our sample consists only of applicants who have been accepted into TFA and applied for TGL funding. Furthermore, only 9% of TGL applicants in our experiment decline TFA's offer, though about 25% fail to actually show up to the first day of school. To attenuate such attrition, 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 assigned 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, with one exception: 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. 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. See Figure A-1 in the appendix for average TGL funding packages based on applicant's expected contribution.

IV. ORIGINAL DESIGN AND RESULTS

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, for the control group, a baseline award package was constructed for each applicant using a modified version of the TGL formula. Figure A-1 shows average grant and loan awards.

In the first year, applicants were equally randomized into either a baseline TGL package, a package with an additional $600 in grants, or a package with 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

Our empirical strategy is based on the nature of the results from Coffman et al. (2019), so we summarize those to begin. Figure I shows the estimated causal increase in teaching on the first day of school for high-need applicants based on additional grant or loan money in Coffman et al. (2019) (replicating Figure V).

Three features of these results drive our empirical approach.

First, as can be seen comparing the first two bars in Figure I, 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 grants or loans) affects our outcomes of interest.

Second, the treatment effects are concentrated among the bottom decile in what TFA refers to as "expected contribution". This value—what TFA calculates that applicants will be able to provide themselves—is calculated using detailed financial data that TGL applicants submit to TFA. This "highest-need"

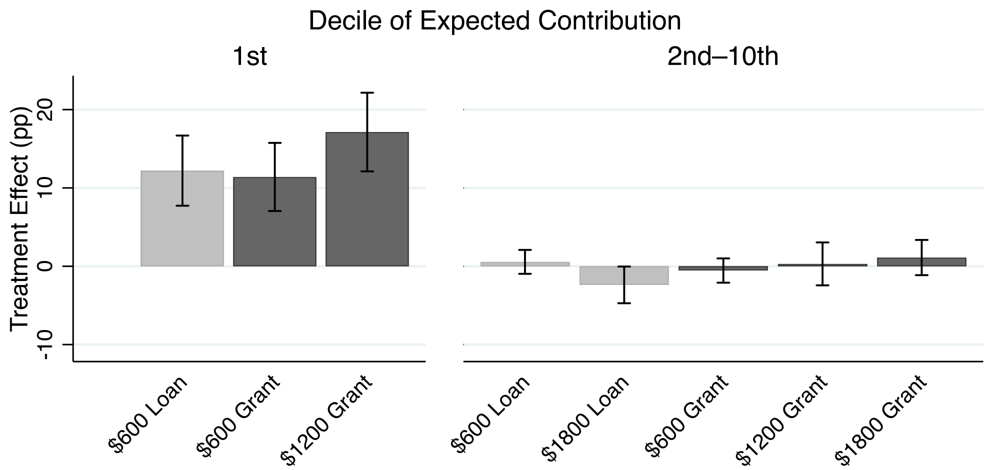

FIGURE I
Treatment Effects on Teaching First Day of School in Coffman et al. (2019)

From Coffman et al. (2019). Figure shows regression-adjusted treatment effects pooled across all years of the experiment. The left set of three bars show the treatment effects observed among applicants in the 1st decile of expected contribution. The right set of bars show the treatment effects observed among applicants in the 2nd–10th deciles of expected contribution. The sample includes applicants from all three years of the experiment (2015–2017). Error bars show standard errors.

decile corresponds almost exactly with the group of applicants whose expected contribution is negative, and as such these applicants are quite financially distressed. The average highest-need applicants has $241 in their checking account, $6,500 of credit card debt, and $19,700 of private student loan debt (which, unlike federal student loan debt, cannot be put into forbearance during their time with TFA).

Third, in addition to finding large effects of additional liquidity for the highest-need group, the experiment in Coffman et al. (2019) found that these effects of liquidity were rather linear (Note the treatment effect of providing an additional $1,200 in grants, the third bar in Figure I, is roughly twice the size of the treatment effects observed from providing an additional $600 of grants or loans, i.e. the first two bars).[3] 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.

Combining these strategies, our main 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.[4] It also includes additional demographic controls in $\mathbf{X}_i$. ε_i is an error term.

[3]Table A-7 shows no evidence of non-linearities in our extended data set.

[4]Multiple times over the course of the year, new applicants were randomized into treatments, so there are many batches in each year of the experiment. For additional details, see Coffman et al. (2019).

V. DATA

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 constructed their TGL packages.

In addition, for each individual whom 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. These administrative data are complete for all individuals in our sample; however, they do not give us an indication of what individuals do if they are not in TFA, including whether they are teaching elsewhere. Both the survey and publicly available data aim to fill these gaps.

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). 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).[5]

Through financial incentives, response rates were high: 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, we complement these 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,

[5]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 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.

graduate school, and graduate school graduation year (if they attended).

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.[6] 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).

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. RAs also recorded the type of school they taught at (if teaching), the industry of their non-teaching job (if not teaching), and any additional schooling and degrees earned. This process yielded at least some employment information on 6,036 of the 7,295 individuals (i.e., 85% of individuals in our sample).

[6]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 and the search process, see Appendix A.I.

VI. RESULTS

In this section, we combine these data sources to analyze the impact of providing additional liquidity on teaching through TFA and teaching anywhere, immediately following the provision of liquidity and up to two and a half years later.

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

Table I displays the effect of every \$100 of additional liquidity on an individual teaching through TFA 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 ($p < 0.001$), a result reported in Coffman et al. (2019). Our new administrative data shows that this 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 even years after the additional liquidity was offered. However, even though the data cannot reject that the later estimated

TABLE I

EFFECTS ON TFA PARTICIPATION

	Replication	New Data			
	Fall Y1	Spring Y1	Fall Y2	Spring Y2	Complete
Extra Liquidity ($100s) × Highest Need	1.80***	1.70***	1.58***	1.54***	1.53***
	(0.41)	(0.44)	(0.46)	(0.46)	(0.47)
Extra Liquidity ($100s) × Not Highest Need	0.01	0.02	-0.04	-0.01	-0.01
	(0.09)	(0.10)	(0.11)	(0.11)	(0.11)
Highest Need	-11.82***	-11.41***	-12.45***	-13.01***	-12.25***
	(2.89)	(2.96)	(3.07)	(3.09)	(3.10)
N	7295	7295	7295	7295	7295
p-value: Highest Need Effect Equal to Fall Y1		0.59	0.44	0.37	0.41
Control Mean: Highest Need	61.1	57.1	50.9	49.1	47.3
Control Mean: Not Highest Need	75.8	72.1	66.6	65.6	63.1

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 are equal to the initial impact (see the p-values in the row labeled "p-value: Highest Need Effect Equal to Fall Y1"), the estimates directionally decrease over time. This decrease could be indicative that teaching is not financially sustainable for a small proportion of those nudged in through our intervention. However, the decrease could also be statistical chance. With our data limitations, we can say if there is a decrease, it seems to be small even after two years of teaching. Finally we observe no effect of liquidity on those who are not in the highest need group, both initially (as in Coffman et al. (2019)) and beyond.

Understanding the lasting impacts of upfront grants and upfront loans separately is also important. If the lasting effects of grants and loans are both positive and similarly sized, we can reasonably conclude that increased liquidity is responsible. Table A-1 shows both grants and loans have substantial, significant positive effects through the completion of the two-year TFA commitment. If anything, upfront loan money has larger (though not statistically different) effects. Either policy is effective at allowing financially constrained candidates to become, and remain, teachers, and the estimates suggest they are equally effective. Liquidity produces TFA teachers, even two years later.[7]

However, as noted above, these results only account for individuals becoming teachers through TFA and only follow individuals for two years. Consequently, in the next section, we report results including our additional outcome data.

[7] Table A-7 reports treatment effects separately for $600 grant, $600 loan, and $1,200 grant bumps, all by need. Though a bit noisy, there do not appear to be any substantial non-linearities for grants and loans of these magnitudes.

VI.B. The Effect of Liquidity on Teaching Anywhere

Next, we analyze the impact of additional liquidity on the likelihood of teaching anywhere, both within TFA or outside of TFA, and during the potential two-year TFA commitment as well as after. Constructing a measure of "teaching anywhere" is simple in most cases. We say someone is teaching in a given academic year if we find affirmative evidence in the TFA data, in the survey responses, or in the publicly available data.[8]

For a few cases, however, constructing a measure of teaching anywhere additionally requires answering two questions: What to do when the data disagree? And what to do when the data are missing? Though infrequent, the data can disagree when the survey data disagree with the publicly available data. Typically, we code them as teaching if either data set shows evidence they are teaching. However, we allow one exception: 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.[9] 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 contradict it.[10]

[8]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.

[9]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.

[10]As discussed later, Table A-4 in the Appendix provides an alternative solution and finds similar results

Finally, how should we handle "missing data", i.e. in cases where, in an academic year, we have no TFA administrative data, no survey data, and no publicly available data for the individual? We present two approaches to these missing cells. In the first approach, we code all missing data as "not teachers". This would likely attenuate any treatment effect, so we consider this likely conservative. Additionally, this was the *ex ante* plan. In the second approach, we drop the missing data from the analysis. This was suggested by the editor and a referee.[11]

Table II shows the results from our main specification for each of the three academic years after an individual is offered funding through TGL, for both approaches to handling missing data. In Year 1, each extra $100 in liquidity offered to high-need individuals by the TGL program increased the likelihood individuals are teaching by 1.17 percentage points (assuming not teaching) or 1.12 percentage points (dropping missing data). These treatment effects represent 62–65% 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.

It is also worth noting that the coefficient on the highest need group for Year 1 is negative, significant and large for either approach, suggesting that

[11]If extra liquidity is correlated with having missing data, this could bias the estimates in either approach. Table A-2 predicts being "found", i.e. have a non-missing value for an academic year. Reassuringly the experimental treatments do not predict missing values across our three data sets.

TABLE II

EFFECTS ON TEACHING ANYWHERE

	Code Missing as Not Teachers			Drop Missing		
	Year 1	Year 2	Year 3	Year 1	Year 2	Year 3
Extra Liquidity ($100s) × Highest Need	1.17***	0.70	0.44	1.12***	0.75**	0.64
	(0.37)	(0.43)	(0.48)	(0.35)	(0.36)	(0.49)
Extra Liquidity ($100s) × Not Highest Need	0.02	0.00	-0.11	0.02	0.05	-0.06
	(0.08)	(0.10)	(0.12)	(0.07)	(0.08)	(0.12)
Highest Need	-7.19***	-5.14*	2.14	-9.96***	-5.29**	-3.85
	(2.61)	(2.83)	(3.14)	(2.46)	(2.42)	(3.18)
N	7295	7295	7295	6633	6181	5749
p-value: Highest Need Effect Equal to Year 1		0.08	0.11		0.51	0.23
Control Mean: Highest Need	73.5	68.6	56.6	78.7	81.6	65.0
Control Mean: Not Highest Need	82.2	76.9	57.9	90.3	89.9	72.9
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. Year 3 estimates do not include TFA data, as they are not available. The first three columns code all missing data as "not teaching" while the last three columns drop missing data from the analysis. 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.

highest need individuals who are not offered extra liquidity are less likely to become teachers. High need individuals with a desire to teach cannot do so at the same rate as their wealthier peers.[12]

The estimates for Year 2 show that liquidity still has an impact into the second academic year. The coefficient estimates are somewhat smaller, at 0.70 and 0.75 percentage points per $100 of liquidity, 60% and 67% of the effect observed in Year 1.

Finally, the Year 3 estimates show that the impact of liquidity is estimated to be positive three years later, although neither the 0.44 or 0.64 treatment effect is statistically significant ($p = 0.35$ and $p = 0.19$ respectively). 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. In fact, neither Year 3 estimate can reject equality with the Year 1 effects (see row labeled "p-value: Highest Need Effect Equal to Year 1").

When interpreting the results from Year 3 (i.e., the last two columns of the table), it is worth emphasizing two points. First, 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.[13] Second, in Year 3, any potential TFA commitment has ended.

[12]The highest need individuals in the control group are also somewhat more likely to be found than their counterparts with less financial need (See Table A-2), possibly because of the higher incentives offered to the highest need individuals to take our 2018 survey.

[13]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

As a result, the baseline rate of teaching drops substantially (See control means at bottom of table). Even if liquidity may have an impact, demand to be a teacher beyond the first two years for this sample has dissipated.

The Appendix provides some robustness checks of the results. To test if the results are being driven by potentially noisy, hard-to-collect data, Table A-3 reruns the analysis recoding data the RAs coded as "low confidence" as missing. Table A-4 revisits how we resolved data disagreements. This analysis codes anyone as a teacher if any data source indicates they are teaching that year, including aspirational data. The estimates in both tables are similar to those above, and typically directionally larger.

The evidence suggests that offering teachers liquidity in the months before they would begin teaching can increase the number of teachers in the short term and maybe even into the medium term.

VI.C. Value of publicly available data: Attenuating sample attrition

Another contribution of this paper is demonstrating the feasibility of using publicly available data to supplement other forms of data. As a demonstration of the value of these data, Appendix Table A-5 replicates the structure of Table II but imagines we did not have access to the publicly available data.

Comparing the results in Appendix Table A-5 to those in Table II, we see that without the publicly available data, the coefficient estimates on extra liquidity, our main variable of interest, differ from the estimates in Table II. Roughly speaking, the public data search found more teachers who were not

near-perfect measures of who is in TFA, so this is (only) concerning for the supplemental data sources.

teaching through TFA. This is consistent with the imbalance in who is "Found" in our non-publicly available data (See the last two columns of Table A-2).

Taken together, the results show 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. For those interested, Appendix A.I details how we collected those data.

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 results suggest targeted upfront liquidity can be a very cost-effective policy to recruit teachers. As a back-of-the-envelope calculation, suppose a policymaker borrowed \$100 at a 6% interest rate and offered an interest-free loan (with the same terms as in our experiment) to a high-need applicant. Using the estimates in Table 2 (thus netting out crowding out teachers from elsewhere), this extra \$100 loan would increase the number of teachers by 0.0117, 0.007, and 0.004 in years one, two, and three. This sums up to a cost of about \$328 per teacher-year. Note that the increased teacher-years beyond year one are not always statistically significant. Even if we conservatively assume that these effects are all zero, the cost of one extra teacher-year increases only to $647, less than 1-1.5% of a teacher salary in most cases. Policymakers should be aware of this potentially cost-effective tool for recruiting new, low-SES teachers.

References

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.

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.

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.

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

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.

Nguyen, Tuan D, Chanh B Lam, and Paul Bruno, "Is there a national teacher shortage? A systematic examination of reports of teacher shortages in the United States," *Annenberg Institute at Brown University working paper series*, 2022.

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.

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, "The state of racial diversity in the educator workforce," 2016.

_ , "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.

Appendix

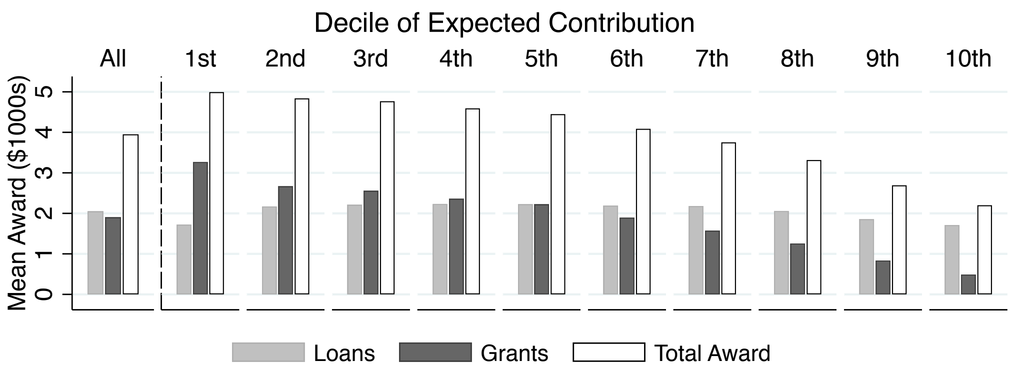

FIGURE A-1

TGL funding in Coffman et al. (2019), by Decile of Expected Contribution

Control awards are the awards that would be offered to applicants randomized into the control group and to which additional funding from the experimental treatments was added. Figure shows the mean loan, mean grant, and mean total control award, both across the entire sample (leftmost group of bars) and broken down by decile of expected contribution (all other groups of bars).

TABLE A-1

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)	(0.44)	(0.46)	(0.47)	(0.47)
Extra Grants ($100s) × Not Highest Need	0.05	0.04	-0.00	0.04	0.04
	(0.10)	(0.11)	(0.12)	(0.12)	(0.12)
Extra Loans ($100s) × Highest Need	2.06***	2.24***	2.04***	1.97***	1.62**
	(0.69)	(0.71)	(0.75)	(0.75)	(0.76)
Extra Loans ($100s) × Not Highest Need	-0.05	-0.02	-0.10	-0.09	-0.08
	(0.11)	(0.12)	(0.13)	(0.13)	(0.13)
Highest Need	-12.15***	-12.07***	-13.01***	-13.56***	-12.39***
	(3.00)	(3.06)	(3.17)	(3.19)	(3.20)
N	7295	7295	7295	7295	7295
p-value: Highest Need Grants Effect Equal to Loans Effect	0.63	0.33	0.44	0.46	0.88
p-value: Highest Need Grants Effect Equal to Fall Y1		0.47	0.40	0.34	0.45
p-value: Highest Need Loans Effect Equal to Fall Y1		0.55	0.96	0.85	0.41
Control Mean: Highest Need	61.1	57.1	50.9	49.1	47.3
Control Mean: Not Highest Need	75.8	72.1	66.6	65.6	63.1

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.

TABLE A-2
PREDICTING BEING FOUND IN PUBLICLY AVAILABLE DATA

	Baseline			Conditional on Teaching Year 1		Only High Confidence			No Public Data	
	Year 1	Year 2	Year 3	Year 2	Year 3	Year 1	Year 2	Year 3	Year 1	Year 3
Extra Liquidity ($100s)	0.15	0.11	-0.12	-0.30	-0.19	0.07	0.06	-0.25	0.83^{**}	-0.44
× Highest Need	(0.20)	(0.35)	(0.33)	(0.29)	(0.35)	(0.21)	(0.36)	(0.34)	(0.32)	(0.47)
Extra Liquidity ($100s)	-0.01	-0.05	-0.08	-0.07	-0.07	-0.01	-0.02	-0.04	-0.05	0.08
× Not Highest Need	(0.06)	(0.08)	(0.09)	(0.06)	(0.09)	(0.06)	(0.08)	(0.09)	(0.08)	(0.11)
Highest Need	2.48	-0.87	7.80^{***}	-0.14	6.34^{***}	2.90^{*}	-0.76	9.18^{***}	-0.82	18.37^{***}
	(1.54)	(2.32)	(2.24)	(1.79)	(2.35)	(1.57)	(2.38)	(2.32)	(2.39)	(3.12)
N	7295	7295	7295	5974	5974	7295	7295	7295	7295	7295
p-value: Highest Need Effect Equal to Year 1		0.89	0.45		0.54		0.95	0.37		0.01
Control Mean: Highest Need	93.4	84.1	87.2	94.6	90.4	93.4	83.2	86.3	81.4	54.9
Control Mean: Not Highest Need	91.1	85.5	79.4	95.3	84.6	90.7	84.9	77.5	82.8	35.0
TFA Admin Data	✓	✓		✓		✓	✓		✓	
Survey Data	✓		✓		✓	✓		✓	✓	✓
Public Data	✓	✓	✓	✓	✓	✓	✓	✓		

Notes. Table shows Linear Probability Model (OLS) regressions of whether an applicant's employment data were found in publicly available data by our RAs for any given academic year. Robust standard errors are reported in parentheses. *, **, *** denote $p < 0.10$, 0.05, and 0.01, respectively. "Baseline" columns use all data. "Conditional on Teaching in Year 1" only uses Year 2 and Year 3 data for those whose employment data was found in Year 1. "Only High Confidence" only utilizes data that the RA coders rated as high confidence (a rating of 2 or 3 on a 0-3 scale), recoding low confidence data as missing. "No Public Data" only uses TFA administrative data and survey 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
USING ONLY HIGHER-CONFIDENCE PUBLIC DATA

	Code Missing as Not Teachers			Drop Missing		
	Year 1	Year 2	Year 3	Year 1	Year 2	Year 3
Extra Liquidity ($100s) × Highest Need	1.22***	0.72*	0.41	1.24***	0.85**	0.71
	(0.37)	(0.43)	(0.48)	(0.35)	(0.35)	(0.49)
Extra Liquidity ($100s) × Not Highest Need	0.02	0.00	-0.07	0.02	0.03	-0.06
	(0.08)	(0.10)	(0.12)	(0.07)	(0.08)	(0.12)
Highest Need	-7.45***	-5.50*	2.70	-10.71***	-5.94**	-4.40
	(2.62)	(2.84)	(3.16)	(2.46)	(2.38)	(3.21)
N	7295	7295	7295	6612	6142	5630
p-value: Highest Need Effect Equal to Year 1		0.07	0.08		0.55	0.17
Control Mean: Highest Need	73.0	68.1	55.8	78.2	81.9	64.6
Control Mean: Not Highest Need	82.2	76.8	56.4	90.5	90.4	72.8
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. In contrast to Table II, 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). The first three columns code all missing data as "not teaching" while the last three columns drop missing data from the analysis. Year 3 estimates do not include TFA data, as they are 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-4
EFFECTS ON TEACHING ANYWHERE
PRIVILEGING ASPIRATIONAL SURVEY DATA

	Code Missing as Not Teachers	Drop Missing
	Year 3	Year 3
Extra Liquidity ($100s)	0.60	0.80*
× Highest Need	(0.47)	(0.46)
Extra Liquidity ($100s)	-0.15	-0.11
× Not Highest Need	(0.12)	(0.12)
Highest Need	2.48	-3.46
	(3.14)	(3.04)
N	7295	5749
Control Mean: Highest Need	58.8	67.5
Control Mean: Not Highest Need	59.0	74.3
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 II, 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-5
EFFECTS ON TEACHING ANYWHERE
EXCLUDING PUBLIC DATA

	Code Missing as Not Teachers			Drop Missing	
	Year 1	Year 2	Year 3	Year 1	Year 3
Extra Liquidity ($100s) × Highest Need	1.63***	1.58***	0.57	1.05***	1.54**
	(0.40)	(0.46)	(0.45)	(0.34)	(0.66)
Extra Liquidity ($100s) × Not Highest Need	-0.06	-0.04	0.03	-0.03	0.04
	(0.09)	(0.11)	(0.10)	(0.06)	(0.19)
Highest Need	-9.99***	-12.45***	7.38**	-11.12***	-7.92*
	(2.82)	(3.07)	(2.89)	(2.46)	(4.51)
N	7295	7295	7295	6040	2718
p-value: Highest Need Effect Equal to Year 1		0.87	0.03		0.70
Control Mean: Highest Need	65.0	50.9	28.3	79.9	51.6
Control Mean: Not Highest Need	78.0	66.6	21.6	94.3	61.8
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.

TABLE A-6

EFFECTS ON TEACHING ANYWHERE
CONDITIONAL ON TEACHING IN YEAR 1

	Code Missing as Not Teachers		Drop Missing	
	Year 2	Year 3	Year 2	Year 3
Extra Liquidity ($100s)	-0.25	-0.21	0.04	-0.04
× Highest Need	(0.36)	(0.51)	(0.25)	(0.49)
Extra Liquidity ($100s)	-0.07	-0.18	0.00	-0.16
× Not Highest Need	(0.07)	(0.12)	(0.05)	(0.11)
Highest Need	-1.00	5.04	-0.96	-0.15
	(2.29)	(3.37)	(1.63)	(3.22)
N	5974	5974	5640	5006
Control Mean: Highest Need	89.16	69.88	94.27	77.33
Control Mean: Not Highest Need	92.1	68.1	96.6	80.5
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.

TABLE A-7
EFFECTS ON TFA PARTICIPATION
FOR GRANTS AND LOANS SEPARATELY
SPLITTING $600 AND $1200 TREATMENTS

	Replication	New Data			
	Fall Y1	Spring Y1	Fall Y2	Spring Y2	Complete
$600 Extra Grants	12.67***	11.94***	9.24**	8.55*	8.65*
× Highest Need	(4.25)	(4.37)	(4.55)	(4.59)	(4.61)
$600 Extra Grants	-0.49	-0.73	0.01	0.11	-0.63
× Not Highest Need	(1.51)	(1.57)	(1.64)	(1.65)	(1.68)
$1200 Extra Grants	20.32***	18.71***	18.47***	18.07***	18.68***
× Highest Need	(4.86)	(5.34)	(5.67)	(5.74)	(5.80)
$1200 Extra Grants	-0.59	-2.00	-1.27	-1.63	-2.23
× Not Highest Need	(2.61)	(2.83)	(3.02)	(3.07)	(3.14)
$600 Extra Loans	13.05***	14.15***	12.20***	11.64**	9.52**
× Highest Need	(4.38)	(4.47)	(4.66)	(4.69)	(4.75)
$600 Extra Loans	0.38	-0.48	-1.17	-1.64	-2.38
× Not Highest Need	(1.48)	(1.55)	(1.63)	(1.64)	(1.67)
Highest Need	-12.93***	-13.26***	-13.08***	-13.75***	-12.96***
	(3.30)	(3.34)	(3.43)	(3.46)	(3.46)
N	7295	7295	7295	7295	7295
p-value: Highest Need $1200 Grants Effect Equal to 2x $600 Effect	0.50	0.52	1.00	0.91	0.87
p-value: Highest Need $600 Grants Effect Equal to Fall Y1		0.69	0.21	0.15	0.20
p-value: Highest Need $1200 Grants Effect Equal to Fall Y1		0.48	0.60	0.54	0.69
p-value: Highest Need $600 Loans Effect Equal to Fall Y1		0.56	0.77	0.64	0.28
Control Mean: Highest Need	61.1	57.1	50.9	49.1	47.3
Control Mean: Not Highest Need	75.8	72.1	66.6	65.6	63.1

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.

FIGURE A-2
Histogram of RAs' confidence ratings
of each coded entry in public data

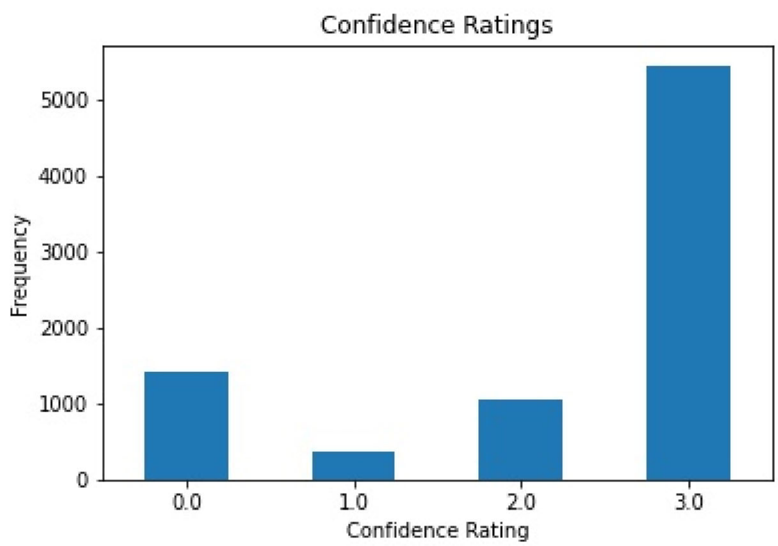

ONLINE APPENDIX

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.