NBER WORKING PAPER SERIES

THE EFFECTS OF LETTERS OF RECOMMENDATION IN THE YOUTH LABOR MARKET

Sara B. Heller
Judd B. Kessler

Working Paper 29579
http://www.nber.org/papers/w29579

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2021

This work was funded by the Social Policy Research Initiative at J-PAL North America. We thank the New York City Department of Youth and Community Development, the New York State Department of Labor, and the New York City Department of Education for sharing the data with us. Thanks to Charlie Brown, Alicia Modestino, Mike Mueller-Smith, Alex Rees-Jones, Ana Reynoso, Basit Zafar, and the JIM group at Princeton for helpful comments. We are particularly grateful to Julia Breitman at DYCD for all of her help along the way, to Ben Cosman at DOE for all his support, and to Alex Hirsch and Lauren Shaw for phenomenal project management and research assistance. All views in the paper are those of the authors and do not represent the positions of any data provider or government agency. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

The Effects of Letters of Recommendation in the Youth Labor Market
Sara B. Heller and Judd B. Kessler
NBER Working Paper No. 29579
December 2021
JEL No. C93,I21,J2,J48

ABSTRACT

Youth employment has been near historic lows in recent years, and racial gaps persist. This paper tests whether information frictions limit young people's labor market success with a field experiment involving over 43,000 youth in New York City. We build software that allows employers to quickly and easily produce letters of recommendation for randomly selected youth who worked under their supervision during a summer youth employment program. We then send these letters to nearly 9,000 youth over two years. Being sent a letter generates a 3 percentage point (4.5 percent) increase in employment the following year, with both employment and earnings increases persisting over the two-year follow-up period. By posting our own job advertisement, we document that while treatment youth do use the letters in applications, there is no evidence of other supply-side responses (i.e., no increased job search, motivation, or confidence); effects appear to be driven by the demand side. Labor market benefits accrue primarily to racial and ethnic minorities, suggesting frictions may contribute to racial employment gaps. But improved employment may also hamper on-time high school graduation. Additional evidence indicates that letters help improve job match quality. Results suggest that expanding the availability of credible signals about young workers—particularly for those not on the margin of graduating high school—could improve the efficiency of the youth labor market.

Sara B. Heller
University of Michigan
Department of Economics
611 Tappan Street, Lorch Hall Room 238
Ann Arbor, MI 48109
and NBER
sbheller@umich.edu

Judd B. Kessler
The Wharton School
University of Pennsylvania
3620 Locust Walk
Philadelphia, PA 19104
and NBER
judd.kessler@wharton.upenn.edu

An appendix is available at http://www.nber.org/data-appendix/w29579

1 Introduction

Youth employment lags farther behind adult employment than can be explained by school engagement, and it recovers more slowly in the wake of major shocks. Even a decade after the Great Recession, youth employment rates in summer—when teenagers are most likely work—were still hovering near their sixty-year low (DeSilver 2021). And the Covid-19 pandemic disproportionately harmed youth labor market prospects at the outset (Inanc 2020; Kochhar and Barroso 2020). Employment gaps are particularly bleak for minority youth. Unemployment and disconnection rates are 30–80 percent higher for Black and Hispanic youth than for their White peers (Bureau of Labor Statistics 2020, 2021; Spievack and Sick 2019), a pattern of disproportionate labor market involvement that is not new (Sum et al. 2014). That youth, and minority youth in particular, have such difficulty securing work is troubling given the range of evidence suggesting that labor force attachment during adolescence and young adulthood may shape employment and wage trajectories for decades (Baum and Ruhm 2016; Kahn 2010; Neumark 2002; Oreopoulos, Wachter, and Heisz 2012).

Half a century of active labor market programs have tried to address these issues, often with relatively costly efforts to improve human capital or provide job search assistance. Despite some success in developing countries and in U.S. sector-focused training (Card, Kluve, and Weber 2018; Katz et al. 2020; Crépon and Van Den Berg 2016), the frequency with which training programs fail to improve youth employment in the U.S. raises the possibility that important frictions may limit young people's access to the labor market. Their short or nonexistent work histories may leave little way to credibly signal future productivity. Employers may engage in statistical or animus-driven discrimination based on age, class, or race, or they may view participation in training programs as a negative signal itself. Alternatively, on the supply side, young people may lack the social networks, knowledge, or confidence to convert their experience into better employment outcomes.

In this paper, we explore the potential role of information frictions in the youth labor market by testing an intervention designed to mitigate them. We partner with the New York City Summer Youth Employment Program (NYC SYEP), which employs city youth to work over the summer, to run a large-scale field experiment with over 43,000 youth. We provide a random subset of youth participants with letters of recommendation from their SYEP supervisors. To make letter production on this scale feasible, we invite program supervisors to complete a survey tool, developed by our research team, that automatically turns their survey responses on individual participants into full-text letters of recommendation. When supervisors agree to produce a letter of recommendation and provide high enough ratings of a worker, that treatment youth receives a digital copy and five hard copies of the letter.

We follow treatment and control youth in administrative data for at least two years after letter distribution to measure their labor market and educational outcomes. Across a pilot after the summer of 2016 and a full-scale study after the summer of 2017, a total of 43,409 SYEP participants are in our main study sample.[1] We measure employment and earnings using unemployment insurance data from the New York State Department of Labor. We observe educational outcomes, which could be directly affected by the letters or indirectly via changes in labor force involvement, using data from the NYC Department of Education.

Partnering with the NYC SYEP provides an ideal environment to assess the role of frictions in the youth labor market. SYEPs are large social programs that provide paid work to youth—often low-income and minority youth—during the summer months. For about half of these youth, SYEP participation is their first experience in the labor market. Consequently, SYEP participants are representative of the groups likely to face barriers in their attempts to capitalize on early work experience. Indeed, while prior literature has found that SYEPs improve important outcomes including criminal justice involvement and mortality, multiple randomized controlled trials suggest they do not have consistently positive average effects on future employment (Davis and Heller 2020; Gelber, Isen, and Kessler 2016; Modestino 2019).

Our results suggest a sizable impact of the letter of recommendation intervention. We find that being sent a letter increases the likelihood that a young person is employed by over 3 percentage points in the year after receiving the letter, a 4.5 percent increase relative to the 70 percent of their control group counterparts who work.[2] Employment effects persist over time, with impacts remaining positive and statistically significant over the cumulative two-year follow-up period; youth who are sent a letter are 2 percentage points more likely to have a job over the next two years, a 2.3 percent increase relative to the control complier mean of 84 percentage points. Cumulative earnings are at least 4.4 percent higher for those sent the letter ($p=0.10$), with different adjustments for skewness suggesting considerably bigger and more statistically significant effects (between 10 and 20 percent, $p<0.05$).

That simply providing a few pieces of paper improves employment and earnings suggests an important role for information frictions in the youth labor market. Our treatment could be mitigating frictions on either the supply side or the demand side. On the supply side, letters may give youth information about what makes them valuable to employers and the confidence to apply to jobs; on the demand side, letters may give employers a clearer signal

1. Our empirical strategy involves stacking panels for the two cohorts, so youth can appear in the data more than once. In total, we have 43,409 observations on 41,633 unique individuals.

2. This effect is 250% as large as prior estimates of the effect of the summer program itself on employment. Gelber, Isen, and Kessler (2016) finds that the NYC SYEP increased employment by 1.2 percentage points in the post-program year by encouraging youth to participate in SYEP again.

about the abilities of a particular youth or make a youth's application more salient.

To assess the mechanism driving our results, we ran an additional data-collection exercise to measure job-seeking behavior among a subset of our sample. We invited 4,000 participants from both treatment and control groups in the 2017 cohort to apply for a short-term online job working for us. Youth in our treatment group were no more likely to apply for our job and no more likely to check a box asking to be considered for a more-selective, higher-paying opportunity, suggesting that our employment effects are not being driven by increased motivation, job search, or confidence.[3] That there was no detectable difference in application behavior among treatment and control youth suggests that the letters work by changing how employers view applicants, rather than how applicants behave.

The only behavioral difference was in the use of the letter itself: Treatment youth were 267% more likely to submit a letter of recommendation as part of their application (4.5 percent of control applicants and 16.5 percent of treatment applicants included a letter in their application). Given the lack of other supply-side responses, it seems possible that this difference in letter usage is the key driver of outcome changes, i.e., that the letters work *only* when employers actually see them. If so, we might view the treatment-control difference in letter use as an implied first stage for the letters' effects on compliers who use the letter in job applications. Back-of-the-envelope calculations—which involve scaling our employment estimates by the implied first stage to approximate the relevant LATE—suggest that actually using the letter is associated with up to a 15 percent increase in employment in the first year and about $1,400 in additional earnings (also about 15 percent) over two years.

In the presence of demand-side frictions, finding additional ways to communicate credible information about youth applicants to potential employers may help youth succeed in the labor market. Consistent with this idea, we find that employment and earnings increases are concentrated among youth who are more highly rated by their SYEP employers, suggesting an important role for information transmission in the letters. We also find that the labor market benefits accrue primarily to minority youth, despite the fact that they receive letters with lower average ratings than White youth. This result suggests that minority youth may face larger frictions in the youth labor market.

One might worry that letters could lead employers to *incorrectly* update beliefs about candidates. If employers' priors about who obtains letters lead them to believe applicants will be higher productivity employees than they actually are, short-term increases in employment could represent bad matches and be followed by increased churn. Because we can track how

3. While it is possible that the increased outside employment among treatment youth affected the decision to apply to our job, application rates did not differ either between treatment and control youth who were employed elsewhere during the quarter we solicited applications, or between treatment and control youth who were not employed elsewhere.

many consecutive quarters employees work at the same employer, we can test this question directly. We find no evidence of increased turnover; treatment youth work more quarters at the same number of jobs, such that some job spells get longer. This set of findings suggests the letters are actually helping workers and employers make more successful job matches.

We collect information on educational outcomes in addition to employment data for two reasons. First, letters could have a direct educational effect if shown to teachers or guidance counselors: other work has shown that teachers' and other adults' beliefs about young people directly affect their outcomes, even when the information that changed those beliefs is fictitious (Rosenthal and Jacobson 1968; Bertrand and Duflo 2017). Letters could also help with college applications if young people have few other sources of recommendations. Second, working during high school could pull young people out of school. There is a general consensus in the current literature that working a small amount has little effect on schooling but that working more than 20 hours is harmful (Buscha et al. 2012; Staff, Schulenberg, and Bachman 2010; Monahan, Lee, and Steinberg 2011; Baum and Ruhm 2016; Ruhm 1997). However, the lack of exogenous variation in this literature means that it is still unclear whether a shock to employment would have a causal effect on school success, at least outside of a setting that mandates continued school enrollment as part of offering a term-time job (Le Barbanchon, Ubfal, and Araya 2020).

For the nearly 20,000 youth in our study who we can observe in New York City public high schools, we find few significant changes in educational performance. But among the subset of youth for whom we can observe graduation, we find that letters of recommendation slow down—but do not appear to stop—high school graduation by pulling people into the labor market. For the average high school student, the welfare implications of efforts to reduce information frictions in the labor market depends on how future career trajectories are affected by the value of additional work experience, and how those benefits compare to the cost associated with a longer time spent in high school. In the meantime, given the heterogeneity we document, policy efforts to provide employers with credible information on minority high school students may minimize the risk of substitution away from school by targeting those not on the margin of graduating on time.

While our employment effects are large, they are consistent with past research that suggests providing even a small amount of information about job-seekers can be quite powerful. In response to fictitious applications in audit studies (Agan and Starr 2018; Kaas and Manger 2012) and to the suppression of information in the labor market (Bartik and Nelson 2019; Doleac and Hansen 2020), employers show less discrimination when they have more information about candidates. Providing performance information has also been shown to increase short-term employment in two different kinds of labor markets. Pallais (2014) randomly hires nearly 1,000 workers on an online labor market platform (oDesk), provides randomly selected workers with more-detailed public performance reviews, and finds that workers with no prior experience benefit on the platform from being hired and rated, while those with prior work experience benefit from the detailed reviews over the next two months. Abel, Burger, and Piraino (2020) find that encouraging a subset of 1,300 job seekers in South Africa to secure letters of recommendation increases job search and survey measures of being employed among women, but not overall, after three months.

We build on this prior work by exploring the impact of letters of recommendation among 43,000 young people in a large, urban U.S. labor market. The setting expands the study of information frictions to an environment where, unlike oDesk or South Africa, employers can potentially access a range of richer signals about youth applicants (e.g., more widespread, visible employment histories or knowledge about local high schools and GPAs) but also face higher hiring costs (e.g., a minimum wage or more burdensome paperwork). Using administrative data, we can observe employment and earnings at jobs across New York State for two years after treatment, as well possible spillovers on high school outcomes for the study youth who are also enrolled in New York City high schools. Our data also allow us to explore specific questions unanswered by the prior literature, such as whether short-term increases in employment are simply a (temporary) result of encouraging bad matches.

Our study provides new evidence that information frictions do prevent young people, especially non-Whites, from securing successful employment. Given that employers seem to value credible information about applicants, finding additional ways to provide personalized information could help improve labor market outcomes among low-income, minority populations like those in our study. Such interventions might be best targeted at those who are not on the margin of graduating on time to avoid harming educational attainment.

The impact of scaling up efforts to facilitate letters of recommendation (or other credible signals) will depend on general equilibrium effects that we cannot directly measure within our study. Welfare effects of expanding letter distribution in general equilibrium could go in either direction. It is possible that reducing information frictions could increase overall employment by helping employers fill vacancies they would otherwise have left open in the face of too much uncertainty. It is also possible that youth with recommendation letters may simply displace those without them (although this is unlikely to have happened within the context of our control group, given that there are about one million 15- to 24-year-olds in the NYC labor market and we sent fewer than 9,000 letters across two years). Even the welfare implications of full displacement are not obvious, however, since policymakers may value the distributional changes or efficiency gains from better matches, even if there were no net change in employment.

Additional research on exactly how letters change employers' decision-making processes would help to predict the welfare consequences of broader efforts to facilitate credible productivity signals. For now, this study provides new evidence on the role of information frictions in discouraging employment and earnings for young people. Such factors potentially limit the impact of programs designed to improve their skills and future labor market outcomes. Fortunately, it may not be particularly costly to reduce these frictions by communicating to employers about applicants' strengths.

2 Setting, Experiment, and Data

2.1 Setting

We partner with the New York City Summer Youth Employment Program (NYC SYEP), implementing our experiment with youth who participated in the summer of 2016 or the summer of 2017. The NYC SYEP is administered by the NYC Department of Youth and Community Development (DYCD). Since a post-Great Recession minimum enrollment of 29,416 youth, enrollment grew steadily to nearly 70,000 youth in 2017. In our program years, the NYC SYEP provided youth with six weeks of paid work during July and August. All NYC residents aged 14–24 were eligible to apply for the SYEP program, though 40% of eventual participants were aged 16–17. Participants in the program were provided with jobs with private sector (45%), non-profit (41%), and public sector (14%) employers. The NYC SYEP directly pays youth for their work with their matched employers at the New York State minimum wage ($11.00/hour in 2017). Youth payroll totaled $83 million in 2017, or roughly $1,200 per youth participant. The NYC SYEP had a total program cost of $127 million in 2017. Over 80% of this cost was funded by the City of New York, with a majority of its remaining funding coming from New York State (see *SYEP Annual Summary* 2017).

2.2 Letter of Recommendation Experiment

We received SYEP data from DYCD on a subset of participants from the 2016 NYC SYEP (n=16,478) and all of the participants in the 2017 NYC SYEP (n=66,763). The program data identified each youth's summer work site and the supervisor or supervisors for the youth at that work site. Using these data, we limited our sample in several ways. First, since we needed to contact supervisors to ask them to complete the letter of recommendation survey, we excluded youth supervised by someone without an email address in the data. Second, we excluded some youth at large work sites to avoid making the survey unmanageable for a single supervisor. In particular, if any supervisor was linked to more than 30 treatment youth, then we randomly selected 30 treatment youth to be included in the survey. We applied the same restriction for the control youth in the survey.[4] In total, this left a sample of 69,222 SYEP participants who were included on at least one survey. Figure 1 traces through this and the subsequent steps of how youth moved through the study.

To generate recommendation letters, we built a survey tool that sent a personalized survey to each supervisor asking about the youth who they supervised that summer (i.e., the youth linked to them in the DYCD data).[5] The email inviting each supervisor to participate explained the letter of recommendation program, included a link to the personalized survey tool, and encouraged them to participate (a sample of the email from 2017 is shown as Appendix Figure A.1). Supervisors were given approximately two weeks to complete the survey, and we sent up to two reminder emails to supervisors who had not yet completed it. For the 2016 cohort, we emailed 3,297 supervisors at the end of September (initial emails went out on 09/29/16). For the 2017 cohort, we emailed 11,877 supervisors in October (initial emails went out on 10/12/17).

The survey began with a brief explanation for supervisors that if they rated a youth positively enough, their responses to the survey questions might be used to construct letters of recommendation. A link to an example letter was provided to aid in the explanation. Respondents were then asked to confirm that they had been a SYEP supervisor during the preceding summer (see screens at the start of the survey in Appendix Figure A.2). Once a respondent confirmed being a supervisor, they were shown the list of treatment youth linked to them in the DYCD data, listed alphabetically by last name.[6] Supervisors selected which youth they had directly supervised and were asked a set of questions about each youth (supervisors were asked about the youth they selected in random order). The survey asked supervisors for an overall rating on the youth's performance and whether they would be willing to answer questions that would turn into a letter of recommendation for the youth (see Figure A.2 for screenshots of the survey). If they were willing, they were also invited

4. To ensure that neither the treatment nor control group exceeded the 30-person-per-survey limit, we randomly assigned treatment and control status prior to making these sample restrictions. Since youth were randomly selected to be excluded, random assignment is still only a function of random variables.

5. The data did not link every youth to a single supervisor. Sometimes, multiple supervisors were listed for a single work site, such that it was not clear which youth reported to which supervisor or if a youth reported to multiple supervisors; in these cases, we assumed the latter for the purposes of constructing our survey tool. Consequently, youth could be listed on more than one survey. Sometimes, a single supervisor was listed for multiple work sites. If the names of the work sites suggested they might be connected (e.g., multiple branches of the same store), we treated them as one work site for the purposes of constructing the survey tool. In the survey, we asked supervisors to confirm the youth that worked for them and to provide the names of others who might have supervised youth so we could include them in the letter of recommendation program as well. If more than one supervisor rated a young person, we generated the letter from the survey with the highest rating, breaking ties by prioritizing letters that included employer contact information, and then those with the most positive responses about the youth.

6. Note that confirming one's identity and position as an SYEP supervisor is how we count "starting" the survey, a definition that is relevant below.

to include their contact information on the letter of recommendation to serve as a reference (97 percent of eventual letters included contact information). They then rated the youth on several attributes, shown in Figure 2.

After the supervisors answered questions about treatment youth, they were asked one question each about control youth—the same question about the overall rating on the youth's performance—all on one screen (see Appendix Figure A.3). They were told that these youth would not be included in the letter of recommendation program.

A total of 5,854 supervisors (39 percent of all supervisors we emailed) opened the survey and confirmed that they had supervised SYEP youth during the preceding summer. In total, 43,409 young people were on a started survey, 29,887 (69 percent) of whom were given an overall rating by employers.

The software we built for this project converted the supervisors' survey responses on treatment youth into formatted letters of recommendation populated with sentences for each youth attribute. For each positively rated attribute, the letter included a dynamically constructed sentence. For example, if in response to the question "How was $<$ *youth name* $>$ at communicating?" the supervisor selected "Very effective," a sentence would appear in the letter that read: "$<$ *Youth name* $>$ was a very effective communicator." Whereas, if the supervisor selected "Not effective" or "Somewhat effective" in response to that question, the sentence about communication would not be included in the letter.

We assigned each attribute to a potential paragraph. If the supervisor rated the youth positively enough on enough attributes to construct a particular paragraph, the paragraph was included in the letter. As long as two paragraphs could be included, the letter was generated for the youth. This procedure ensured that any letters of recommendation our survey tool generated had enough positive things to say about the youth to provide a positive letter that would not be too sparse. Our software produced letters of recommendation as PDFs on official DYCD letterhead. The letters ended with "Sincerely," followed by the name of the supervisor and work site. A short note in the footer of the letter described our letter of recommendation pilot program. Figure 3 shows a sample letter.

In total, we generated and sent 8,780 letters (1,805 in 2016 and 6,975 in 2017). We uploaded digital copies of these letters to Dropbox with a link sent to the youth for whom emails were known (1,737 in 2016 and 6,720 in 2017).[7] In addition, we mailed five physical copies of the letters via USPS to each youth along with a cover letter providing context and suggested uses for the letter (see Appendix Figure A.4 for a sample cover letter; similar

7. About 56 percent of letter recipients clicked the link in their email to view the letter digitally. Many SYEP youth create an email solely for the purpose of the online application and then abandon it, so some letter recipients may not have seen the email containing the link to the digital copy of the letter.

text was sent to youth via email along with the link to the soft copy of the letter).[8] All letters of recommendation were sent in time for winter holiday hiring in the year after SYEP participation (letters were sent to youth in early-December 2016 for the 2016 cohort and in mid-November 2017 for the 2017 cohort).

2.3 Job Application Data

To understand the mechanisms through which letters of recommendation might impact labor market outcomes of treatment youth, we advertised a job to a subset of the youth in our data, solicited job applications, and hired youth ourselves. We composed a job listing for a one-time, remote, paid work assignment, emailed the job listing to 4,000 randomly selected subjects from our 2017 cohort, and observed their job application behavior. The sample was evenly split among treatment and control youth from the letter of recommendation experiment (i.e., youth who had been eligible and ineligible to receive the letter of recommendation) who also had an email address in the data so we could send them the job application.

The job was described as being with a professor at the University of Pennsylvania who was looking for former NYC summer job participants for a short-term and flexible job. The job description highlighted several qualifications: "responsible," "self-motivated," having an "enthusiastic approach," and offered compensation of $15/hour. A link to an application with a deadline to submit an application was included at the bottom of the job description (see the email invitation sent to youth with the job description in Appendix Figure A.5).

Youth who clicked the link in the email were taken to a job application that asked a few standard contact, background, and employment experience questions. Our application also provided an optional space to upload up to three "supporting documents (e.g. resume or other documents that might strengthen your application)." The application did not explicitly mention uploading letters of recommendation, but it would have been easy for youth to upload the soft copy of the letter of recommendation provided to them in our experiment (see the screenshot of this prompt in Appendix Figure A.6).[9] This upload interface allowed us to measure whether youth provided supporting materials—including a letter of recommendation—with their applications and to assess whether this differed by treatment and control youth.

Finally, to assess the confidence of youth in our study, we gave applicants the opportunity

8. Of the 8,780 sets of letters mailed to youth, 127 were returned as undeliverable.

9. We intentionally avoided explicitly mentioning a letter of recommendation to see if youth in our study would choose to upload a letter without a specific prompt to do so. We saw this as realistic to job applications in practice where a youth could choose to provide a potential employer with a letter of recommendation even if one was not specifically requested.

to check a box on the application to be considered for a more selective, higher-paying position ($18/hour) that required a stronger application. The application told them explicitly that being considered for the more selective position would not affect their chances at being selected for the regular job.

All those who submitted an application that included their name, email address, and at least 1 additional field were hired.[10] The job itself was an online survey of multiple choice questions. These questions asked youth about their experiences job-seeking and considering college, as well as about their career and education goals. At the end of the survey, there were free-response questions about the youth's experience in SYEP.[11] Workers were instructed to finish everything they could within a two-hour time frame. All youth who initiated the job-task (n=227) were paid for two hours of work via a mailed, pre-loaded debit card (so our job does not appear in the administrative data on employment and earnings).

2.4 SYEP Administrative Data

Administrative data from the NYC SYEP comes from the NYC DYCD, which runs the program. We received data on a subset of participants of the 2016 NYC SYEP and all participants of the 2017 NYC SYEP. The data on SYEP participants include identifiers (e.g., name, date of birth, and social security number) that allow us to match to various data sources; demographics (e.g., gender, race, and pre-SYEP education status) that allow us to test for balance across treatment and control; and contact information (e.g., mailing address and email address) that we used to send letters of recommendation to treatment youth. We define racial/ethnic categories based on the self-reported categories in the application, making the classifications mutually exclusive (e.g., "White" only captures non-Hispanic Whites). We also received information on the work site where the youth worked for the summer and information about the supervisors at that work site, including name and email address. We use the information on work site and supervisor to send the letter of recommendation surveys.

2.5 NYS Department of Labor Data

We obtained earnings and employment data from the New York State Department of Labor (NYSDOL). Data came from NYSDOL's quarterly Unemployment Insurance (UI) dataset, which covers formal sector employment, excluding self-employment or farming income. The

10. To ensure our hiring for the more selective job was incentive compatible with our instructions about higher selectivity, the youth needed to click the box asking to be considered and needed to complete one or more of the open response questions in addition to fulfilling the requirements for the standard job.

11. Youth hired for the more selective job were asked additional open response questions that required more thoughtful consideration.

data include employer name, employer FEIN, employer address, employer NAICS, and amount paid in each quarter. NYSDOL analysts matched SYEP participants to UI data using social security number. When multiple profiles in the NYSDOL data shared the same social security number, we used name to disambiguate the UI data. In total, 99 percent of SYEP youth in our letter of recommendation experiment were matched to the NYSDOL data with no difference between treatment and control youth ($\beta = 0.001, p = 0.128$).[12]

We use data from Q1 (January–March) of 2010 through Q4 (October–December) of 2019. This window provided considerable baseline data as well as two years of outcome data after letters were sent to SYEP participants in our treatment group for each study cohort.[13]

2.6 NYC Department of Education Data

Education data come from the NYC Department of Education (DOE).[14] The DOE used name, date of birth, and gender to perform a probabilistic match between our study sample and their records between the 2015–2016 and 2019–2020 school years, inclusive. SYEP applicants fail to match because they never appear in the DOE system (e.g., always attended private school), matched to more than one student record (DOE treats multiple matches on the same name and birth date as a non-match), or because typographical errors or name changes prevented identifying a study participant's education records. Overall, 88 percent of our main sample matched to a DOE student record, with no treatment-control difference in match rates ($\beta = -0.003, p = 0.359$). Within the sample that matched to a DOE student record, 7,643 had no active enrollment within our 2015–2020 data. These students were largely old enough to have left school prior to 2015 (their average age at randomization is 19.7), although some may have transferred to private or non-NYC districts prior to the start of our data. This leaves 69.9% of our main sample with at least some education information in the data, with no treatment-control difference ($\beta = -0.003, p = 0.436$).

12. In theory, everyone in our data should have matched to the data, since they were all listed as a SYEP participant during the summer prior to the program. Some of the non-workers may not have matched to the UI data despite having worked due to typographical mistakes or incorrect SSNs. Others may not have ever been paid by SYEP despite being listed as a participant in their data, and so not actually have received any wages to be reported to the UI system.

13. Letters were sent in Q4 (October–December) of 2016 or 2017, depending on cohort. Consequently, we have additional quarters of data for the youth in the 2016 cohort, but we limit the analysis to the period we can observe for full years for both cohorts, and we stop prior to any influence from Covid-19.

14. At the request of the data provider, when we merge DOE data with the rest of our study data, we exclude the self-reported citizenship status that appears on the SYEP application, so that education outcomes are never linked to citizenship status. SYEP application data also provides spotty information on whether youth live in public housing or are on public assistance; those are also never linked to DOE data.

3 Method of Analysis

This section discusses how we perform the analysis in this paper. In Section 3.1, we describe our sample definitions and our outcomes of interest for each data source. In Section 3.2, we describe our empirical approach, including our regression specifications. In all sections, we note cases where we deviated from our pre-analysis plan with accompanying explanations for these choices.[15]

3.1 Sample Definitions and Outcomes

3.1.1 Labor Market Sample

Our main sample to explore labor market outcomes consists of the 43,409 SYEP participants who were on a survey that a SYEP supervisor started (i.e., the SYEP participant appeared on at least one survey in which the supervisor clicked the link inviting them to take the survey and confirmed on the first page of the survey—prior to viewing which youth were on the survey or what their treatment status was—that they supervised youth that summer). This excludes the 25,813 youth who were randomized and placed on a survey that no supervisor ever opened.

We pre-specified this subsample of youth on a started survey as a key sample of interest, because neither treatment nor control youth on *unopened* surveys could have actually received treatment. This kind of non-compliance mechanically reduces statistical power and is orthogonal to treatment status, so we focus on the subsample with a first stage of 0.404 (rather than the first stage of 0.254 when we include youth on unopened surveys).[16] As a result, the treatment effect of receiving a letter of recommendation in our main sample is representative of the population of youth whose supervisors both had an up-to-date email address in the DYCD data and were willing to click on an invitation to participate in the letter of recommendation program. The estimates from this sample of youth might differ from the treatment effect on the broader sample of all SYEP youth, because different types of youth are placed into jobs with different types of supervisors.[17]

15. The pre-analysis plan can be found at https://osf.io/8zwdr/

16. While we pre-specified this subsample as a key sample of interest, our main sample included all SYEP participants that we randomized, because we did not anticipate that only 39% of supervisors would open the survey and that such a large fraction (i.e., over one third) of the sample would be on an unopened survey. For completeness, we present and discuss results for this larger sample in Appendix Section A.8. We choose to emphasize the results from our smaller sample in the main text, because the power gains from focusing on this subsample give better insight into the effect of the letter of recommendations on the sample of youth who might actually have been eligible to be treated, given the actions of their supervisors.

17. Appendix Section A.8 shows that youth who were on unopened surveys are indeed observably different than the youth in our control group of opened surveys on demographics and employment outcomes, although not in their likelihood of applying to our job posting. As such, it is possible that forcing supervisors to rate

Since supervisor non-response was driven by an inability to reach supervisors by email or by a lack of supervisor interest or capacity to complete the survey, limiting our analysis to this sample does not interfere with the integrity of random assignment (i.e., until the supervisors reached the substantive survey questions, they had no way of knowing which youth would be included in the survey or which youth would be in the treatment or control groups). As discussed below, Table 1 shows that our main sample is balanced across treatment and control youth.

3.1.2 Labor Market Outcomes

We pre-specified a primary focus on annual earnings, winsorized to deal with outliers, along with alternative methods to adjust for skewness as robustness checks. We pre-specified an indicator for any employment as a secondary outcome. Our main analysis shows employment and earnings in the first year after randomization (4 quarters including the quarter the letters were sent), the second year, and cumulatively, winsorized at the 99.5^{th} percentile, as well as log(earnings+1). Since the +1 transformation effectively manufactures the proportional change from zero, in Appendix Section A.1.1 we also show robustness to alternative transformations (winsorization at the 99^{th} percentile, adding 0.1, 10, and 100 to earnings prior to logging, and the inverse hyperbolic sine).

We also pre-specified exploratory analyses on: (1) the number of jobs and length of jobs to assess job stability and match quality, and (2) the industry of employment to assess whether letters help youth find jobs in which they now have experience (i.e., those over-represented in SYEP jobs) or whether the letters help market youths' skills to the higher-paying industries that are under-represented in SYEP jobs (see a discussion of these industry definitions in Gelber, Isen, and Kessler (2016)). For (1), we define a job spell as all consecutive quarters worked at the same employer. Other outcomes related to spell length and industry are discussed in Appendix A.1.2 and A.1.3.

3.1.3 Job Application Sample

We randomly selected 2,000 control youth and 2,000 treatment youth from our main 2017 cohort to invite to apply to our job application.[18] Table A.5 shows baseline balance for this subsample. Although the vast majority of baseline covariates are balanced, we note that the treatment group in this subsample is significantly more Hispanic by chance (33 percent in the treatment group versus 29 percent in the control group). As we show in Appendix Section A.6, labor market impacts for Hispanic youth are larger than for other groups. As

youth might have somewhat different effects than those we estimate here.

18. We also invited 1,000 youth from unopened surveys (i.e., outside of our main sample) to ensure that job application behavior was not dramatically different for the youth excluded from our main sample.

a result, the point estimates for employment and earnings are considerably larger for this sample, although the smaller sample size makes the estimates imprecise (see discussion in Appendix Section A.2).

3.1.4 Job Application Outcomes

For the sub-sample of individuals we randomly selected to receive our job application advertisement, we pre-specified three key outcomes: whether someone applied, whether they uploaded a letter, and whether they checked the box to apply to a more selective job as a measure of confidence.

Observing whether there is a treatment-control difference in application rates helps us to test whether there is a supply-side job search response behind any potential changes in labor market outcomes. The proclivity to opt into consideration for the "more selective" job tests for treatment-control differences in self-efficacy and motivation or in beliefs about the probability of being hired. Whether applicants uploaded a letter provides a measure of how much letter use actually changed in job applications.

We also report two additional outcomes to provide a more complete picture of job application behavior: whether someone clicked the link to view the job application (regardless of whether they applied), and whether someone uploaded any file (e.g., CV, transcript, or anything else) in support of their application.

3.1.5 Rated Youth Sample

To distinguish whether the letters provide a general signal or convey useful information about worker productivity, we report labor market impacts separately for youth who receive low overall ratings (categories 1–4: "Very Poor," " Poor," "Neutral," and "Good") versus high overall ratings (categories 5–7: "Very Good," "Excellent," and "Exceptional") from supervisors on the question about overall performance, asked of both treatment and control youth. Unlike our main sample, however, there is the potential for selection into who receives a rating based on supervisor behavior in the survey. Because the survey was designed to maximize the number of letters generated, treatment youth were listed first, along with a longer, multi-page set of questions on each youth; control youth were all listed at the end of the survey on a single page, with check boxes that allowed the supervisor to quickly answer the single overall quality question about each control youth. The different positioning and survey content for treatment and control youth could change the probability a supervisor rated a particular youth. Additionally, supervisors were told (and could decide whether) their responses would be turned into a letter for treatment youth, but not for control youth. The possibility of sending a letter may itself lead supervisors to make different decisions about whether to rate a youth or which rating to give. Because of both differences, we would not necessarily expect the distribution of treatment and control youth to be identical conditional on having a rating or receiving a particular rating.

In fact, treatment youth are significantly less likely to have received a rating than control youth (66 versus 71 percent, p<.01). Although the distribution of baseline characteristics is not statistically different across groups conditional on having a rating, it is significantly different for youth receiving a low rating (Table A.6 shows that treatment youth receiving a low rating are observably different from control youth receiving a low rating, perhaps because supervisors were more hesitant to give low ratings when a letter might be produced than when they knew it would not).

To minimize the role of selection introduced by whether a youth is rated, when we explore treatment effects by ratings, we focus on the sub-sample of youth who were on a survey in which the supervisor rated every treatment youth and every control youth in the survey. There are 13,911 youth who were on such a survey. Since everyone is rated, these surveys leave no room for treatment-control differences in who is rated within the survey. In this group, treatment youth are only 1 percentage point less likely to appear on a completed survey overall (31.6 versus 32.5 percent, p=0.066), with observables jointly balanced (p=0.865). Appendix Table A.8 shows the distribution of baseline characteristics is also similar within rating groups for this sub-sample. Appendix Section A.3 shows that even without this sample restriction, results are very similar when using all youth with a rating.

3.1.6 Education Sample

Because we knew much less about what education data would be available to us at the time of pre-specification, the education analysis is where we deviate most from our pre-analysis plan. As reported above, about 70 percent of our sample has any active record in the DOE records during the period our data cover (2015–2020). But in practice, many of these students either graduated or left school prior to our 2016 and 2017 study years. And while charter school students do appear in DOE data as having active records, DOE does not share any information about school engagement, performance, or graduation for charter school students with outside researchers. Because of the amount of missing data, including on individual elements like GPA and college enrollment even within individuals that have some educational records, we leave the analysis of separate educational outcomes to Appendix Section A.4. That analysis focuses on students who were enrolled in grades 8–12 in the year prior to randomization, were not enrolled in a charter school at the end of the pre-randomization year, and who had not yet graduated. This is the sample who we expect to be in a DOE high school in the year after SYEP (if they do not transfer or stop attending school).

In the main text, we focus on the most substantively important high school outcome, and the one that provides the greatest contribution to the question of how work matters for schooling: high school graduation. We note that because this focus was not pre-specified, it should be considered exploratory. Graduation data are not available for everyone. Per state standards, DOE only reports graduation in the academic years that correspond to a student's on-time (4th), 5th, or 6th year graduation cohort (even if a student returns to school after their 6th year). Graduation data are missing for students who transfer to a charter school; move out of district; fall under another exclusion, such as having an individualized education plan (IEP); or were not in a 4th–6th year graduating cohort between fall 2015 summer 2020. To avoid conditioning the graduation sample on what could potentially be an outcome (e.g., transferring in or out of the District), we restrict the sample to those most likely to be observed in the graduation data based only on pre-randomization characteristics.

In particular, the main text limits attention to students who were enrolled in grades 10–12 in the year prior to randomization, were not enrolled in a charter school at the end of the pre-randomization year, and who had not yet graduated. This is the part of the high school sample with the most available graduation information prior to the end of the data: within this group, 64 percent of students are old enough to have complete graduation data, and all others have reached their 5th-year graduation date by the end of the data.[19] This sample excludes students outside of the DOE, pre-randomization dropouts and graduates, students who temporarily stopped attending public school or had not yet joined the school district in the year before randomization, and those too young to observe their full graduation data (8th and 9th graders in both study cohorts). Appendix Section A.5 reports results including the 8th and 9th graders for whom we can at least observe on-time (but not later) graduation before the end of the data, and Appendix Figure A.8 diagrams the available graduation data by grade and study cohort.

Our main 10th–12th grade graduation sample contains 13,732 students, with no treatment-control difference on being in this sample either overall or conditional on being in our main education data ($\beta = -0.0006, p = 0.926$). We note that there is some chance imbalance on observables within the education data, discussed in more detail in Appendix A.4.3. One benefit of the post-double-selection LASSO is that it adjusts for chance imbalance in a principled way.

19. We assess graduation cohorts based on the grade a student was in during the pre-randomization year, since we only observe the state-defined cohort that officially determines graduation years if someone appears in the graduation data—a potentially selected group. Within our main graduation sample, only 10th graders in the 2017 cohort have incomplete graduation information (n=4,984 have not yet reached their 6th-year graduation date).

3.1.7 Education Outcomes

We explore three main outcomes in our graduation sample. The first measures whether youth graduated on-time (i.e., within four years of starting 9th grade) at any time during our data. The second measures whether someone graduated within six years of starting 9th grade. The third measures "school persistence," which captures whether someone has either graduated or is still attending school in the last academic year of our data (2019–2020)

For our measures, we count any graduation outcome between the start of the academic year when randomization occurred and the end of our graduation data. This will capture most, but not all, eventual high school graduation. There are two reasons why we miss some eventual high school graduation. The first is that, for the 4,984 youth who are in 10th grade prior to the 2017 summer, our data end before their 6th-year graduation date. The second is that graduation after the 6th year is not recorded in our DOE data, so we will not observe any eventual graduation of students who spend more than 6 years in high school. Our school persistence measure is designed to include these youth, who are still pursuing a diploma.[20]

In the main text, we focus on the relationship between educational attainment and labor force involvement. To measure this relationship, we define a set of mutually exclusive joint outcome indicators: working and graduating, never working but graduating, working and not graduating, and never working and not graduating. We define these indicators for all three of our education attainment measures: on-time graduation, ever graduating, and graduating or still attending school. Note that these joint outcomes measure employment over a 2-year follow-up, while graduation includes either a 3- or 4-year follow-up, depending on the study cohort. The pattern across these outcomes will allow us to assess whether any potential shifts in educational attainment occur among the same group that experiences shifts in employment. Results on the three educational attainment indicators separately, as well as other high school performance measures and on-time college enrollment including the full sample of 8th–12th graders, are in Appendix Section A.4.

20. Sixty-four youth in our graduation sample do not have any graduation information available, likely because they transferred out of the district (or joined a different group excluded from state graduation counts) after randomization. Since these individuals did not receive a diploma from NYC DOE, we assign them zeros for graduation. DOE discharge codes suggest there is no treatment effect on whether students transfer out of the district ($\beta = 0.003$, $p = 0.322$, with a control mean of 0.033). Since we do not observe graduation outside the district, the balance on transfers helps to rule out the possibility of differential mobility biasing the graduation results. In theory, differential mobility could also be an issue for our labor market results, since we only observe UI data within New York state. Although the available post-secondary data is limited to a subset of the full sample, the on-time college enrollment measure discussed in the appendix can help assess whether out-of-state mobility is different across treatment groups. That measure records if someone is enrolled in an out-of-state college 6 months after their on-time graduation date, and it suggests that treatment youth are no more likely to leave New York State for college ($\beta = -0.002$, $p = 0.692$, with a control mean of 0.064).

3.2 Analytical Method

3.2.1 Main Analysis

We begin with an intent-to-treat (ITT) analysis by regressing each outcome variable on a treatment indicator and baseline covariates:

$$Y_{it} = \alpha + \beta T_i + \gamma X_{it-1} + \epsilon_{it}$$

where Y_{it} is an outcome for individual i at time t, T_i is an indicator for random assignment to treatment, and X_{it-1} is a vector of covariates measured at or before the time of random assignment. As pre-specified, we use a post-double-selection LASSO to select which covariates to include in each regression (Belloni, Chernozhukov, and Hansen 2014a, 2014b; Belloni et al. 2012).[21] We always include an indicator variable for study cohort, since randomization occurred separately by study year. Because 1,776 individuals appear more than once in the data, we cluster our standard errors on individual as identified by SSN in the SYEP data.

Not every treatment youth on a started survey was sent a letter, either because: they were on a survey answered by someone who was not their direct supervisor, the supervisor did not want to provide a letter, or the supervisor provided ratings that were not positive enough for a letter to be sent. As a result of this kind of non-compliance, the ITT will understate the effect of being sent a letter. In addition, we cannot observe who actually views or uses the recommendation letters in practice. Instead, we do two things to provide a sense of the effect's magnitude for those who are actually treated. First, we use random assignment as an instrument for whether a youth was sent a letter. Since we perfectly observe whether every youth was sent a researcher-generated letter, we can estimate this treatment-on-the-treated effect for everyone. We report control complier means as a baseline measure to assess proportional changes for compliers (Kling, Liebman, and Katz 2007).

Second, we use our job application data to estimate the proportion of treatment youth who actually use the letter in practice. Because we find evidence that letters do not generate a supply-side response, it is possible that the letters work *only* when youth actually show them to employers. In this case, we can approximate the treatment-control difference in letter use with our job application data and use that as an implied first stage to scale the ITT effect. Because our job posting is not representative of all job applications, this extrapolation involves a strong assumption that letter use in the rest of the labor market looks like letter use in our job application. This assumption could fail in two ways: either

21. We implement this with the Stata commands pdslasso and ivlasso (Ahrens, Hansen, and Schaffer 2020). See Appendix Section A.9 for a list of the covariates we offer the LASSO, and for results without any covariates or with all covariates as robustness checks.

because it is easier to remember the letter or submit the letter in our application than in other real-world applications (in which case we would likely understate this LATE effect), or because treatment changes the composition of who applies to our job posting by changing whether youth are employed when we send our advertisement (the direction of which depends on employment treatment effects). We discuss the interpretation issues further below, but in general we consider this a rough approximation of the effect of actually using the letter for those who choose to use it, not a direct estimate of the relevant LATE.

3.2.2 Heterogeneity analysis

Although we pre-specified at the outset that we would not have enough statistical power to differentiate heterogeneous treatment effects, we follow our pre-analysis plan in conducting exploratory analyses based on the characteristics most likely to affect youth's labor market prospects. For all heterogeneity tests, we report the ITT, the first stage, and the IV separately for each group to show how much of ITT differences are from different rates of receiving a letter and how much are from different responses conditional on being sent a letter.

The main text focuses on separating labor market impacts for White and minority youth, where minority is defined as any non-White self-classification, including Black, Hispanic, Asian, and Other (including American Indian, Pacific Islander, mixed race, or unspecified other). To help identify whether employers are responding to the specific information in the letters, we also test for heterogeneity across supervisor ratings. If the letters are successfully communicating specific information, we would expect that providing letters with higher ratings would generate larger labor market benefits.

Appendix Section A.6 breaks down effects by specific racial and ethnic groups, and it reports heterogeneity on the other pre-specified categories: age, gender, school enrollment (as self-reported on the application to SYEP), and neighborhood. The appendix also explores heterogeneity by previous work experience to see if the effects are limited to those who lack other signals of an ability to get a job on their applications.

4 Results

4.1 Summary Statistics

Table 1 shows average pre-randomization characteristics for the treatment and control groups. No more differences are significantly different than would be expected by chance, nor are they jointly significantly different ($p = 0.267$). Study participants reflect the population that participates in NYC's SYEP. On average, they are just over 17 years old, about 43 percent male, largely identify as minorities (only 12.5 percent list being White on their application), and 75 percent report being in high school in the spring prior to the SYEP. About 45 percent of participants never appear in the UI data prior to their participation in SYEP, but 97 percent work during the SYEP year, earning an average of just under $2,400 that year.

4.2 Labor Market Effects

Table 2 reports the main labor market effects. Panel A shows that being assigned to the treatment group increases employment rates by 1.3 percentage points (1.8 percent relative to the control mean of 70 percent) during the year following letter distribution.[22] Actually being sent the letter increases year 1 employment by 3.1 percentage points (4.5 percent relative to the control complier mean). The point estimates in the second year after letter distribution are still positive but smaller and not statistically significant. However, the increase in employment is still significant over the cumulative two-year follow-up: being sent a letter increases net two-year employment rates by 2 percentage points (2.3 percent).

It is likely that the employment effect will fade out eventually, since almost all control youth will eventually work in the formal labor market at least once. But earnings changes would not necessarily fade out if the letter is helping set youth on a better employment trajectory or find better jobs. Panels B and C report program effects on winsorized earnings in dollars and log(earnings+1), respectively. The treatment effect grows in levels over the two years observed, though effects are somewhat noisy. Those sent a letter earned a total of $433 more over the two-year follow-up period ($p = 0.101$), a 4.4 percent increase. The $log+1$ transformation is more precise, with Panel C showing a significant 18.6 percent increase in cumulative earnings and significant increases in both years 1 and 2.

Because there is a treatment effect on the extensive margin, the results may be sensitive to how we handle the proportional change at $0. Table A.1 shows alternative level, log, and asinh transformations. The results suggest that the magnitude of the change is somewhat sensitive to functional form, ranging up to 23 percent, but generally statistically significant. We focus on the 4.4 percent estimate in the main text both because our pre-specified primary outcome was winsorized earnings and because it is the most conservative estimate. Since we also pre-specified that we would use a range of robustness checks to adjust for skewness, we conclude that the evidence suggests that the letters of recommendation generated a sizable increase in earnings.

Table 3 digs more deeply into the UI records to understand how labor market outcomes are improving. The first column shows that treatment youth work in 0.05 more quarters (0.11 for letter recipients) than their control counterparts. The last column shows that conditional

22. Letters were sent in December of 2016 and November of 2017, so we include the final quarter of each calendar year as the first quarter of year 1.

on working at all (i.e., for those selected into work), treatment youth find jobs sooner than control youth (0.12 quarters sooner for letter recipients). Together, this pattern suggests that the letters help youth shorten the job search process, but do not merely substitute early work for later work; youth work more than they would have otherwise.

One concern about this pattern is the possibility that supervisors could be over-updating their beliefs about youth, interpreting the letters as a stronger quality signal than they actually are. If so, we might expect increased churn, with treatment youth getting hired and fired more frequently than controls. The rest of Table 3 suggests this is not the case: there is no increase in the number of job spells treatment youth have. The point estimate on the number of jobs (including 0s) is positive but not statistically significant, partly capturing the change at the extensive margin. Conditional on working at all (column 3), which is a selected sample, the point estimate is a precise zero. In other words, there is no evidence of additional churn among those who work. And as we would expect for young people who start working earlier and work more in the same number of jobs, Appendix Section A.1.2 shows that some job spells get longer. This provides further evidence that letters are not reducing—and in fact may be increasing—the quality of job matches.

4.3 Mechanisms

4.3.1 Assessing Changes in Labor Supply

A key question about the observed increase in labor market success among treatment youth is whether the letters increase labor supply by increasing youth job search intensity or confidence, or whether the letters increase labor demand by changing beliefs about applicants with letters or increasing the salience of those applicants among employers. By distributing our own job posting to 4,000 treatment and control youth, we are able to generate some evidence on why the letters increase employment and earnings and to approximate how treatment changes letter use in job applications. Appendix Section A.2 shows that we have baseline balance within this sample and shows the main employment results for this group.

Table 4 suggests that supply-side responses (increased job search, motivation, or confidence) are unlikely to be driving the labor market improvements. We find no evidence that treatment youth are more likely to click on the application link or actually apply to our posting.[23] The second column shows that 8.8 percent of the control group and 8.2 percent of the treatment group applied to our job, a difference that is not statistically significant.

23. The "applied" variable here measures whether a youth entered enough information in the application for us to know who filled out the application form. We define "applied" this way because we hired people even if they did not answer all the questions on the application. To actually be hired, the youth additionally needed to click submit on the final page of the application. There is also no treatment-control difference on whether youth were hired per this definition.

We also find no evidence that the letter increased confidence among applicants conditional on applying; treatment youth are no more likely to volunteer for the more selective job than control youth (see the last column of Table 4, which, adjusting for application rates, translates into 60 percent of control applicants and 51 percent of treatment applicants checking the box to apply for the more selective job).

Of course, it is possible that even though the letters did not change the rate at which young people applied to our job, they could have changed the composition of who applied. Since treatment youth were more likely to be employed in the formal labor market, their interest in our short-term, online job may have been directly affected by treatment (even though we framed the job as flexible enough to be compatible with other work). That said, we cannot reject the null that observables are jointly unrelated to treatment status among applicants, suggesting this is not likely to be the case.[24] Additionally, even if we condition on not being employed elsewhere during the quarter the job application was distributed (a selected group), there is still no significant difference in application rates or our confidence measure ($\beta = -0.01$, $p = 0.351$ for applying and $\beta = -0.01$, $p = 0.132$ for checking the selective box).[25] So the lack of an increase in supply-side job-seeking behavior does not appear to be due to treatment youth being more likely to be employed already. Overall, the evidence from our job application suggests that labor market improvements are coming from employers responding to letters of recommendation, not from changes in youth's application behavior or confidence.

4.3.2 Assessing the First Stage

As an important check on whether treatment youth actually use the letters we send them—a necessary condition for employers to be able to respond to the letters—the final two columns of Table 4 show treatment effects on the files job applicants uploaded in their application to our job posting. There is no detectable change by treatment in the probability that youth upload some form of supporting material. But there is a dramatic change in whether youth upload a letter of recommendation. Only 0.4 percent of the control group submits a letter, including zeros for those who do not apply (conditional on applying, this translates to 4.5 percent of control applicants submitting a non-intervention letter with their application). Treatment youth are two and a half times more likely to submit a letter of recommendation than the control group: 1.4 percent of all those invited to apply submit a letter (16.5 percent

24. We test for differences between the treatment and control individuals who applied for our job, conditional on being in our application sample, by interacting each baseline covariate with an indicator for whether the individual applied, regressing treatment on all covariates and these interactions, and then testing the hypothesis that all interaction coefficients are jointly 0. The p-value of this test is 0.14.

25. The same is true conditional on being employed in that quarter: $\beta = 0.0008$, $p = 0.959$ for applying.

conditional on applying). Since about 40 percent of treatment youth actually received a letter, this implies that about 41 percent of letter recipients use them when they apply to a job (16.5 percent relative to 40 percent).

Given the lack of supply side response, it seems reasonable to suppose that letters might only work when employers see them. If so, the observed rates at which letters are used can also benchmark the first stage of letter use, which under the exclusion restriction we can use to extrapolate how big employment responses are for youth who actually use their letters. If we make the quite strong assumptions that the difference in letter use we observe in our job application applies to the entire sample, that treatment and control youth apply to jobs at the same rate, and that everyone applies to at least one job, then the implied first stage for letter use is a 12 percentage point increase (4.5 versus 16.5 percent among applicants). Scaling our main ITT effects by this first stage would in turn imply that the employment increase for those who use the letter is about 15 percent relative to baseline in the first year, and 8 percent over two years, with an additional $1,400 in earnings over that time.

Because of all the extrapolation involved in this calculation, we view it as a back-of-the-envelope estimate. If we think that it might be easier to use the letter in our application than in typical job applications (e.g., because receiving a job advertisement that references SYEP and having a screen to upload supporting material reminds treatment youth about the letter or primes them to use it more than in a typical job application), or that not everyone applies to at least one job, then we are likely overstating the number of treatment youth who used a letter relative to controls. In that case, our 12 percentage point first stage would be an overestimate, and the actual LATE would be even larger than our calculations here suggest.

4.3.3 Assessing Changes in Labor Demand

The evidence so far suggests that employers are the ones responding to the increased use of letters of recommendation in the job applications of treated youth. A final mechanism question is how those letters affect their hiring behavior. It is possible that because letters are infrequently included in typical job applications, it is the presence of the letter itself—regardless of content—that makes an application more salient, resolves some basic uncertainty about whether an applicant is likely to show up at all, or overcomes statistical discrimination. Alternatively, employers may be using the content of the letter to try to discern something more nuanced about future employee reliability or productivity.

Although we did not send letters where SYEP supervisors included too few positive comments about the youth they supervised, there is still variation in how positive supervisors were in their letters that allows us to assess whether employers respond to letter content. Table 5 shows employment and earnings effects separately for youth who received low ratings (categories 1–4, corresponding to “Very Poor,” “ Poor,” “Neutral,” and “Good”) and high ratings (categories 5–7, corresponding to “Very Good,” “Excellent,” and “Exceptional”).[26] Highly rated youth were much more likely to receive a letter (81 percent versus 33 percent). So the ITT differences between the groups reflect both differences in letter receipt and differences in outcomes conditional on being sent a letter, although the substantive pattern of results is relatively similar for both the ITT and TOT.

The first result of note is that the ratings do seem to capture attributes that matter in the labor market. Looking at the control means, low-rated youth are 6 percent less likely to work during year 1 (67 versus 72 percent employed), though they catch up to high-rated youth over time. They also earn just under $1,500 (14.6 percent) less over 2 years. Second, we find that, cumulatively, the low-rated group has net employment effects close to 0 and cumulative earnings point estimates that are negative but with huge standard errors. In contrast, the high-rated group has employment and earnings effects that grow over time, such that they are driving basically all of the net positive effects of the treatment.

It appears, then, that employers are using the substance of the letters to identify those who are likely to be highly productive employees, but who might not otherwise be noticed during the hiring process. While one might wonder whether the low-rated group simply chooses not to use letters in their job applications, results from our job application suggest otherwise (see Appendix Section A.3). For every 100 letters sent to high-rated treatment youth, we received 3 job applications that included letters. For every 100 letters sent to low-rated treatment youth, we received 4 applications including letters. So there is no indication that low-rated letter recipients are less likely to use letters when applying for jobs.[27] The group of young people who did not impress their SYEP supervisors as much may need more intensive investment in improving skills to reap long-term gains.

4.4 Work and Graduation

Education results on engagement and performance outcomes are in Appendix Section A.4. In general, there is little evidence that letters improve student performance in school (e.g., by changing teacher or guidance counselor beliefs or encouraging college application). While none of the treatment effects are statistically significant for the full education sample, there is one pattern that becomes significant in several subgroups and alternative specifications: on-time (4-year) graduation shows a substantively important decline, while the point estimates

26. Note that if youth received an overall rating less than “Good,” the paragraph that included text about the overall rating was not printed in the letter. These letters were still produced, though, as long as enough other attributes were rated positively enough.

27. While high-rated letter recipients apply at somewhat higher rates and use letters somewhat more often, many more of them are sent letters than their low-rated counterparts.

on ever graduating (including delayed graduation), school persistence (graduating or still attending), and enrolling in college immediately after 4-year graduation, are much closer to zero. This pattern provides suggestive evidence that, on average, recommendation letters increase employment while slowing down—but not stopping—high school graduation for those still in school.

A natural hypothesis is that our employment and education effects are driven by the same group of youth: that by pulling young people into the labor force, the letters make it harder for them to complete their high school education on time, leading them to spend additional time in school. Table 6 tests this hypothesis by showing IV effects for the joint outcomes of working and graduating, never working but graduating, working but not graduating, and never working nor graduating. Appendix Table A.13 is the ITT version of this table. Panel A shows these outcomes using on-time graduation; Panel B uses whether someone ever graduated—on time or otherwise—within our data; and Panel C shows school persistence (i.e., whether someone graduated or is still attending school in the final year of our data).

The pattern of results is consistent with letters reducing on-time graduation—but not eventual graduation—for the same group of youth who also have positive employment impacts from the letters. The second column of Panel A shows that treatment significantly decreases the proportion of youth who graduate on time without working by 2.1 percentage points for letter recipients. Treatment also generates a corresponding increase, of roughly similar size (3.0 percentage points), in the proportion of youth who work but do not graduate on time, as shown in column 3. Since there are no significant changes in the other two categories, this is suggestive of a shift from graduating on time without work to working but failing to graduate on time. The additional 0.9 percentage point increase in working without graduating on time appears to come from the group that neither works nor graduates on time, as shown in column 4, suggesting a modest employment increase from youth who would not have graduated on time even if they had not received the letter.

By contrast, Panel B shows that for ever graduating, there is no significant change in the proportion of youth who work but do not graduate. Rather, treatment generates a significant 2.6 percentage point decrease in the proportion of letter recipients who never work but graduate, and an offsetting increase of 2.4 percentage points in the proportion who both work and graduate. In other words, the letters seem to encourage employment among those who graduate. The point estimates among non-graduates show a similar pattern of a negative point estimate for not working with an offsetting positive point estimate for working. These estimates, however, are substantially smaller and not statistically significant.

Panel C measures educational attainment as school persistence: either graduating or continuing to attend school. The basic pattern of results is similar to the results in Panel B.

Among those who do not persist, we see a decline in never working and an offsetting increase in working. We see the same pattern among those who do persist.

Appendix Table A.16 shows that these shifts are concentrated among students who have below-median GPAs in the year before randomization, consistent with the idea that it is students struggling in school who are most responsive to the letter and whose educational attainment is most harmed by the increase in work. Appendix Section A.5 discusses these and additional robustness checks, including showing a similar but slightly noisier set of results when 8th and 9th graders, who have had less time to graduate, are included in the analysis.

In sum, for the subset of study youth for whom we have the most complete graduation data, this joint outcome analysis suggests that the shock to employment generated by the letters slows down graduation, but does not stop it. These results provide a useful addition to the literature on working during school, which typically has been unable to measure on-time graduation and has hit a ceiling effect when analyzing overall graduation (Buscha et al. 2012; Staff, Schulenberg, and Bachman 2010; Monahan, Lee, and Steinberg 2011; Ruhm 1997; Baum and Ruhm 2016). The welfare implications of slowing down graduation—and whether the slightly longer high school career outweighs the benefits of the additional work experience and earnings—depends on how long and by how much the letters affect the trajectory of future longer-term outcomes.

It is worth emphasizing that only about 30 percent of our overall sample is in this graduation analysis, due to the smaller set of youth who are of relevant age and for whom we have education data. The rest of the youth in our study are either too young for us to observe graduation, are out of school already, or are enrolled in schools that are not in our data. Those not in our high school data still have significant increases in year 1 employment as well as much larger point estimates on earnings than the high school sample (see Appendix Section A.4). So, from a policy perspective, it may be feasible to focus on mitigating information frictions that impact youth who are not still in high school, or at least students not on the margin of graduating on-time.

4.5 Heterogeneity

Table 7 shows that, while we do not have the statistical power to differentiate between the two groups, the labor market effects of the letters appear to be driven by racial and ethnic minorities. The letters have no significant effect on White youth, who show negative but imprecise point estimates from the letters. Effects are only positive and statistically different from zero for the non-White (Black, Hispanic, Asian, and Other) youth in the sample.[28]

28. Appendix Section A.6 shows results separately for the individual racial and ethnic groups, as well as for other subgroups of interest.

The first-stage results suggest part of the difference in the ITT effects is that minority youth are much more likely to be sent a letter (42 percent versus 30 percent for minority youth and White youth respectively). This difference may have to do with differential selection into the SYEP in the first place, since only 12.5 percent of our SYEP participants are White, or with differences in the kinds of SYEP supervisors for whom minority and White youth work. But the IV results show that, even conditional on being sent a letter, the point estimates for employment and earnings are much bigger for minority youth than for White youth. One might wonder whether the larger IV effect reflects differences in letter quality; perhaps letters for minority youth matter more because they are stronger letters. However, we observe that the opposite is true: conditional on receiving a letter, White youth receive ratings that are 0.43 points higher (on our 7-point scale) than minority youth, with no significant differences in whether they use the letter on our job application (see Appendix Section A.7 for descriptions of how letters and job application behavior differ by subgroup). Consequently, it seems that the letters sent to minority youth have a particularly powerful effect because of how employers respond to them.

Table 8 shows that the education results in Table 6 are, in fact, driven by minority students. They are driving the declines in on-time graduation, although the proportion who both work and graduate eventually still increases. The shifts in persistence are consistent with the possibility that the letters help both those who would and would not finish school shift into the labor market. But since multiple outcomes move at once, we cannot rule out the possibility that some youth may be pulled out of school by their increased employment. Interestingly, the substantive patterns are largely in the opposite direction among White students, with hints that the letters are helping White students' educational success without increasing employment. But the small number of White youth in the data means the changes are not statistically distinguishable from zero.

5 Discussion

Sending youth a few copies of a letter and an email with a link to that letter increased employment rates by 4 percent—and perhaps as much as 15 percent for those who actually use the letter. These results provide new evidence that there are, in fact, frictions in the labor market for youth, and minority youth in particular, that are relatively low cost to overcome. We do not find differences in job-seeking behavior among treatment youth other than using the letter, suggesting that employers are the ones responding to the additional information contained in the letters. This interpretation is bolstered by results showing that higher performing youth get a larger labor market benefit from the letters.

We also find that recommendation letters lead to a decline in on-time graduation, driven by substitution toward work on the margin. In addition to being important for any future policy decisions about letters of recommendation or other signals about youth, this finding also speaks to the literature on the impact of working during high school. Our letters provide a plausibly random shock to working during high school, which appears to extend the time spent in high school. To assess the welfare implications of this letter-generated substitution, we would need to make some strong assumptions about how long the increase in earnings will last and how that compares to the longer time in high school. It seems likely that the net effect may not be beneficial, especially if any subgroup leaves school entirely. Reducing employment frictions is most likely to have a net benefit for those who have already finished their high school careers, or at least are not currently on the margin of graduating (although we find few labor market benefits among high-achieving students).

Overall, the labor market results indicate that employers respond to credible information about youth, such that finding additional ways to provide them with personalized information about an individuals' strengths could help improve labor market outcomes among low-income, largely minority populations like those in our study. For social programs looking to help youth or other disadvantaged groups capitalize on their training or early work experience, this is an important insight.

A natural question is how broadly this finding applies—whether we are documenting that letters help overcome the particular stigma associated with SYEP participation or that, more broadly, at least some youth unemployment is due to frictions surrounding the availability of information about young applicants. The answer to this question rests partly on whether the employers in our data knew that youth applicants were SYEP participants, which would be necessary for the stigma story. While we cannot observe that directly, we can take a hint from the applications that youth submitted in response to our job advertisement. In those applications, only 22 percent of applicants self-identify as a SYEP participant in either their list of work experience or their résumé. Given that almost 80 percent of job applicants would not appear to employers as SYEP participants—and that the recommendation letters came on letterhead from the agency that runs the program, increasing the salience of the SYEP—it seems plausible that the frictions we document are not specific to SYEP-related beliefs among employers.

Despite the positive effects of the letters on labor market outcomes, our findings do not necessarily imply that policymakers should try to give everyone letters of recommendation. Our estimates are for youth who receive letters when survey responses are voluntary and responses are positive enough. Effects may differ outside of this population.[29] In addition,

29. It is difficult to say from the observable differences in youth across the opened and unopened surveys whether effects would be bigger or smaller if supervisors were forced to fill out the surveys. The unopened

any effort to generate more widespread use of credible signals like letters of recommendation could result in displacement; youth with letters might gain jobs, but at the expense of those who would otherwise have taken those jobs.

Such displacement and general equilibrium effects are worth considering as part of efforts to scale up such programs. There are several conditions under which a scaled-up version could be beneficial, even with considerable displacement. If policymakers valued equity, then transferring job opportunities to those farther down the income distribution or to historically marginalized groups might be socially beneficial. Alternatively, even with no net change in employment, letters could generate efficiency gains by helping employers and employees find better matches. And, if employers end up leaving some vacancies open in the face of too much uncertainty about applicants, as they appear to in an online marketplace (Pallais 2014), a widespread information-sharing intervention might increase overall employment. Finally, there could also be general equilibrium effects on the supply side; if young people understand that they may receive helpful recommendation letters, they may work harder in their jobs, generating additional productivity as well as better letters to which future employers will respond more positively.

Even in partial equilibrium, our experiment establishes that information frictions prevent minority young people from getting jobs they could otherwise succeed in. Further research into the precise way employers update their beliefs or substitute across workers in response to efforts to mitigate these frictions would be a productive next step in assessing the most effective way to leverage our findings into higher youth employment.

surveys contained more White youth, who have smaller labor market effects. But they also had more youth already out of high school, which could diminish graduation crowd-out, and more youth with work experience prior to SYEP, who have directionally larger point estimates, see Appendix Section A.6.

References

Abel, Martin, Rulof Burger, and Patrizio Piraino. 2020. "The Value of Reference Letters: Experimental Evidence from South Africa." *American Economic Journal: Applied Economics* 12, no. 3 (July): 40–71. Accessed April 5, 2021.

Agan, Amanda, and Sonja Starr. 2018. "Ban the Box, Criminal Records, and Racial Discrimination: A Field Experiment*." *The Quarterly Journal of Economics* 133, no. 1 (February 1, 2018): 191–235. Accessed February 22, 2019.

Ahrens, Achim, Christian Hansen, and Mark Schaffer. 2020. *pdslasso and ivlasso: Progams for post-selection and post-regularization OLS or IV estimation and inference.* http://ideas.repec.org/c/boc/bocode/s458459.html.

Bartik, Alexander, and Scott Nelson. 2019. "Deleting a Signal: Evidence from Pre-Employment Credit Checks." *University of Chicago, Becker Friedman Institute for Economics Working Paper,* nos. 2019-137.

Baum, Charles L., and Christopher J. Ruhm. 2016. "The Changing Benefits of Early Work Experience." *Southern Economic Journal* 83, no. 2 (October): 343–363.

Belloni, Alexandre, Daniel Chen, Victor Chernozhukov, and Christian Hansen. 2012. "Sparse Models and Methods for Optimal Instruments With an Application to Eminent Domain." *Econometrica* 80 (6): 2369–2429. ISSN: 1468-0262, accessed October 29, 2021.

Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen. 2014a. "High-Dimensional Methods and Inference on Structural and Treatment Effects." *Journal of Economic Perspectives* 28, no. 2 (May): 29–50. ISSN: 0895-3309, accessed October 29, 2021.

———. 2014b. "Inference on Treatment Effects after Selection among High-Dimensional Controls†." *The Review of Economic Studies* 81, no. 2 (April 1, 2014): 608–650. ISSN: 0034-6527, accessed October 29, 2021.

Bertrand, Marianne, and Esther Duflo. 2017. "Field experiments on discrimination." *Handbook of economic field experiments* 1:309–393.

Bureau of Labor Statistics, U.S. Department of Labor. 2020. *Employment status of the civilian noninstitutional population by age, sex, and race.* Accessed April 5, 2021.

———. 2021. *Employment and Unemployment Among Youth — Summer 2021.* Accessed November 29, 2021.

Buscha, Franz, Arnaud Maurel, Lionel Page, and Stefan Speckesser. 2012. "The Effect of Employment while in High School on Educational Attainment: A Conditional Difference-in-Differences Approach*." *Oxford Bulletin of Economics and Statistics* 74 (3): 380–396.

Card, David, Jochen Kluve, and Andrea Weber. 2018. "What works? A meta analysis of recent active labor market program evaluations." *Journal of the European Economic Association* 16 (3): 894–931.

Crépon, Bruno, and Gerard J Van Den Berg. 2016. "Active labor market policies." *Annual Review of Economics* 8:521–546.

Davis, Jonathan M.V., and Sara B. Heller. 2020. "Rethinking the Benefits of Youth Employment Programs: The Heterogeneous Effects of Summer Jobs." *The Review of Economics and Statistics* 102, no. 4 (October): 664–677.

DeSilver, Drew. 2021. "In the U.S., teen summer jobs aren't what they used to be." Pew Research Center. Accessed April 5, 2021.

Doleac, Jennifer L., and Benjamin Hansen. 2020. "The Unintended Consequences of "Ban the Box": Statistical Discrimination and Employment Outcomes When Criminal Histories Are Hidden." Publisher: The University of Chicago Press, *Journal of Labor Economics* 38, no. 2 (April 1, 2020): 321–374. Accessed September 13, 2021.

Gelber, Alexander, Adam Isen, and Judd B. Kessler. 2016. " The Effects of Youth Employment: Evidence from New York City Lotteries *." *The Quarterly Journal of Economics* 131, no. 1 (February): 423–460.

Inanc, Hande. 2020. *Breaking Down the Numbers: What Does COVID-19 Mean for Youth Unemployment?* Mathematica Policy Research Reports. Mathematica Policy Research. Accessed April 5, 2021.

Kaas, Leo, and Christian Manger. 2012. "Ethnic discrimination in Germany's labour market: A field experiment." *German economic review* 13 (1): 1–20.

Kahn, Lisa B. 2010. "The long-term labor market consequences of graduating from college in a bad economy." *Labour Economics* 17, no. 2 (April): 303–316. Accessed February 22, 2019.

Katz, Lawrence F, Jonathan Roth, Richard Hendra, and Kelsey Schaberg. 2020. *Why Do Sectoral Employment Programs Work? Lessons from WorkAdvance.* Technical report. National Bureau of Economic Research.

Kling, Jeffrey R, Jeffrey B Liebman, and Lawrence F Katz. 2007. “Experimental Analysis of Neighborhood Effects.” *Econometrica* 75, no. 1 (January): 83–119. Accessed February 22, 2019.

Kochhar, Rakesh, and Amanda Barroso. 2020. “Young workers likely to be hard hit as COVID-19 strikes a blow to restaurants and other service sector jobs.” Pew Research Center. Accessed April 5, 2021.

Le Barbanchon, Thomas, Diego Ubfal, and Federico Araya. 2020. “The Effects of Working while in School: Evidence from Uruguayan Lotteries.” *Available at SSRN 3398385.*

Modestino, Alicia Sasser. 2019. “How Do Summer Youth Employment Programs Improve Criminal Justice Outcomes, and for Whom?” *Journal of Policy Analysis and Management* 38, no. 3 (June): 600–628. Accessed September 13, 2021.

Monahan, Kathryn C., Joanna M. Lee, and Laurence Steinberg. 2011. “Revisiting the Impact of Part-Time Work on Adolescent Adjustment: Distinguishing Between Selection and Socialization Using Propensity Score Matching.” *Child Development* 82 (1): 96–112.

Neumark, David. 2002. “Youth Labor Markets in the United States: Shopping Around vs. Staying Put.” *Review of Economics and Statistics* 84, no. 3 (August): 462–482. Accessed February 22, 2019.

Oreopoulos, Philip, Till von Wachter, and Andrew Heisz. 2012. “The Short- and Long-Term Career Effects of Graduating in a Recession.” *American Economic Journal: Applied Economics* 4, no. 1 (January): 1–29. Accessed February 22, 2019.

Pallais, Amanda. 2014. “Inefficient Hiring in Entry-Level Labor Markets.” *American Economic Review* 104, no. 11 (November): 3565–3599. Accessed February 22, 2019.

Rosenthal, Robert, and Lenore Jacobson. 1968. “Pygmalion in the classroom.” *The Urban Review* 3, no. 1 (September): 16–20. Accessed February 22, 2019.

Ruhm, Christopher J. 1997. “Is High School Employment Consumption or Investment?” *Journal of Labor Economics* 15 (4): 735–776.

Spievack, Natalie, and Nathan Sick. 2019. “The Youth Workforce: A Detailed Picture.” Urban Institute, July 22, 2019. Accessed April 5, 2021.

Staff, Jeremy, John E. Schulenberg, and Jerald G. Bachman. 2010. “Adolescent work intensity, school performance, and academic engagement.” *Sociology of Education* 83 (3): 183–200.

Sum, Andrew, Ishwar Khatiwada, Walter McHugh, and Will Kent. 2014. "Deteriorating Labor Market Fortunes for Young Adults." *Challenge* 57, no. 3 (May 1, 2014): 60–83. Accessed April 5, 2021.

SYEP Annual Summary. 2017. Technical report. New York City Department of Youth and Community Development.

Tables and Figures

Table 1: Descriptive Statistics

	Control Mean	Treatment Mean	Test of Difference
N	21,695	21,714	
Age	17.2	17.2	0.641
Male	0.427	0.427	0.991
Black	0.409	0.411	0.805
Hispanic	0.289	0.289	0.944
Asian	0.129	0.130	0.734
White	0.124	0.125	0.756
Other Race	0.049	0.045	0.080
In High School	0.755	0.751	0.339
HS Graduate	0.044	0.042	0.202
In College	0.173	0.180	0.081
Not in UI Data	0.009	0.011	0.128
Never Employed Pre-SYEP	0.450	0.457	0.113
Ever Worked, Year -4	0.153	0.149	0.210
Earnings, Year -4	318	320	0.882
Ever Worked, Year -3	0.267	0.266	0.840
Earnings, Year -3	585	585	1.000
Ever Worked, Year -2	0.437	0.435	0.627
Earnings, Year -2	1072	1050	0.412
Ever Worked, Year -1	0.966	0.966	0.798
Earnings, Year -1	2379	2368	0.722
No Education Match	0.126	0.123	0.359
In HS Sample	0.454	0.454	0.938
Joint F-Test	$F(24, 41632) = 1.16$, $p=.267$		

Notes: N = 43,409. 390 youth missing race/ethnicity and 1 missing education. Test of Difference reports the p-value from a regression of each characteristic on a treatment indicator, controlling for a cohort indicator and using standard errors clustered on individual.

Table 2: Labor Market Effects

Year	1	2	Cumulative
Panel A: Employment			
ITT	0.0127***	0.0058	0.0079**
	(0.0041)	(0.0041)	(0.0034)
CM	0.701	0.72	0.841
Sent Letter (IV)	0.0313***	0.0144	0.0195**
	(0.0102)	(0.0102)	(0.0083)
CCM	0.697	0.728	0.841
Panel B: Earnings, Winsorized at 99.5th Percentile			
ITT	60.03	110.1	168.66
	(45.89)	(73.23)	(106.86)
CM	3579	5964	9543
Sent Letter (IV)	154.11	281.4	433.17
	(113.40)	(180.95)	(264.02)
CCM	3729	6162	9894
Panel C: Log(Earnings + 1)			
ITT	0.095***	0.059*	0.075**
	(0.033)	(0.035)	(0.030)
CM	5.61	6.08	7.33
Sent Letter (IV)	0.234***	0.146*	0.186**
	(0.081)	(0.087)	(0.073)
CCM	5.64	6.18	7.39

Notes: N = 43,409. Winsorization in Panel B recodes each quarter's highest earnings to the 99.5th percentile of all quarterly earnings before summing across years. CM shows control means; CCM shows control complier means. Regressions include baseline covariates. Standard errors clustered on individual are shown in parentheses. * p<0.1, ** p<0.05, *** p<0.01

Table 3: Amount and Timing of Work

	Num Quarters Worked	Num of Job Spells	Num of Job Spells if >0	Time to First Qtr Worked
ITT	0.045**	0.019	0.002	-0.048**
	(0.022)	(0.014)	(0.014)	(0.020)
CM	3.46	1.98	2.36	2.19
Sent Letter (IV)	0.111**	0.046	0.006	-0.119**
	(0.054)	(0.034)	(0.034)	(0.048)
CCM	3.59	1.98	2.35	2.18
N	43409	43409	36647	36647

Notes: Spells are defined as consecutive quarters with earnings from same employer. Time to First Qtr Worked conditions on having at least one spell. CM shows control means; CCM shows control complier means. Regressions include baseline covariates. Standard errors clustered on individual are shown in parentheses. * p<0.1, ** p<0.05, *** p<0.01

Table 4: Job Application Effects

	Clicked Link	Applied	Checked Selective Job Box	Uploaded Any File	Included Letter of Rec
ITT	-0.007	-0.006	-0.01	0.003	0.010***
	(0.009)	(0.009)	(0.007)	(0.007)	(0.003)
CM	0.103	0.088	0.053	0.052	0.004
Sent Letter (IV)	-0.02	-0.019	-0.027	0.006	0.024***
	(0.024)	(0.022)	(0.017)	(0.018)	(0.007)
CCM	0.138	0.123	0.082	0.065	0.009

Notes: N = 4,000. CM shows control means; CCM shows control complier means. Regressions include baseline covariates. Standard errors clustered on individual are shown in parentheses. * p<0.1, ** p<0.05, *** p<0.01

Table 5: Labor Market Effects for Youth with High and Low Supervisor Ratings

	First Stage	Employment Y1	Employment Y2	Employment Cumulative	Earnings Y1	Earnings Y2	Earnings Cumulative
ITT, Low Ratings		0.0250*	-0.0146	0.002	63.24	-163.28	-98.75
		(0.0133)	(0.0134)	(0.0110)	(129.42)	(211.89)	(308.35)
ITT, High Ratings		0.013	0.0238***	0.0174**	106.42	338.72**	437.84*
		(0.0087)	(0.0087)	(0.0070)	(99.38)	(164.96)	(237.14)
P-value, test of diff.		0.455	0.016	0.235	0.791	0.062	0.168
CM, Low		0.673	0.721	0.836	3109	5409	8518
CM, High		0.715	0.720	0.846	3729	6251	9979
IV, Low Ratings	0.3301***	0.0756*	-0.0442	0.0056	190.62	-506.09	-311.39
	(0.0103)	(0.0405)	(0.0405)	(0.0332)	(391.67)	(641.82)	(933.60)
IV, High Ratings	0.8108***	0.0161	0.0293***	0.0212**	130.92	413.65**	535.92*
	(0.0057)	(0.0108)	(0.0108)	(0.0087)	(122.57)	(203.44)	(292.50)
P-value, test of diff.	0.000	0.155	0.08	0.649	0.884	0.172	0.387
CCM, Low		0.61	0.756	0.821	3012	5942	8950
CCM, High		0.713	0.717	0.843	3626	6079	9714

Notes: To avoid selection into who is rated within a survey, sample includes only youth who were on a survey where the supervisor rated all listed youth (N = 13,911). Low Ratings includes rating categories 1–4; High Ratings includes rating categories 5–7. P-value, test of diff shows p-values for tests of the null hypothesis that treatment effects are equal in low and high ratings groups. CM shows control means; CCM shows control complier means. Regressions include baseline covariates. Standard errors clustered on individual are shown in parentheses. * p<0.1, ** p<0.05, *** p<0.01

Table 6: IV Effects on Joint Employment and Graduation Outcomes

	Panel A: On-Time Graduation			
	Ever Work, On-time Grad	Never Work, On-time Grad	Ever Work, Not On-time	Never Work, Not On-time
Sent Letter (IV)	0.0031	-0.0209*	0.0296***	-0.0105
	(0.0146)	(0.0117)	(0.0103)	(0.0071)
CCM	0.727	0.12	0.114	0.038
	Panel B: Any Graduation			
	Ever Work, Graduated	Never Work, Graduated	Ever Work, Not Graduated	Never Work, Not Graduated
Sent Letter (IV)	0.0238	-0.0257**	0.0065	-0.0048
	(0.0145)	(0.0120)	(0.0092)	(0.0065)
CCM	0.758	0.127	0.086	0.029
	Panel C: Any Graduation or Continued Attendance			
	Ever Work, Persisted	Never Work, Persisted	Ever Work, Not Persisted	Never Work, Not Persisted
Sent Letter (IV)	0.0154	-0.0198	0.0153*	-0.0108*
	(0.0146)	(0.0123)	(0.0086)	(0.0059)
CCM	0.793	0.129	0.051	0.027

Notes: N=13,732. Analysis is conducted on the main graduation sample (non-charter 10th–12th graders in the pre-randomization year, see text for details). First stage for this sub-sample is 0.44. Panel A shows whether someone ever worked during the two-year follow up and whether they graduated on-time (i.e., 4th-year graduation). Panel B shows whether someone ever worked during the two-year follow up and whether they ever graduated (i.e., 4th-, 5th-, or 6th-year graduation). Panel C shows whether someone ever worked during the two-year follow up and whether they either graduated or had positive days attended in the last year of our data. CCM shows control complier means, which may not total to 1 across categories due to estimation error in the IV and the inclusion of different sets of covariates in the post double-selection LASSO. Regressions include baseline covariates. Standard errors clustered on individual are shown in parentheses. * $p<0.1$, ** $p<0.05$, *** $p<0.01$

Table 7: Labor Market Effects for Minority and White Youth

	First Stage	Employment Y1	Employment Y2	Employment Cumulative	Earnings Y1	Earnings Y2	Earnings Cumulative
ITT, Minority		0.0134***	0.0066	0.0090**	79.03	149.27*	227.66**
		(0.0044)	(0.0044)	(0.0036)	(48.13)	(77.52)	(112.60)
ITT, White		0.0048	-0.0019	-0.0031	-70.27	-162.15	-230.22
		(0.0114)	(0.0122)	(0.0096)	(144.77)	(218.80)	(328.57)
P-value, test of diff.		0.483	0.513	0.236	0.328	0.18	0.187
CM, Minority		0.6932	0.7229	0.839	3540	5958	9498
CM, White		0.7518	0.6949	0.851	3754	5702	9457
IV, Minority	0.4188***	0.0319***	0.0158	0.0214**	194.27*	365.42**	557.35**
	(0.0036)	(0.0106)	(0.0105)	(0.0086)	(114.85)	(184.98)	(268.66)
IV, White	0.2973***	0.0157	-0.0077	-0.0112	-241.61	-563.96	-798.88
	(0.0088)	(0.0385)	(0.0412)	(0.0323)	(488.27)	(737.83)	(1107.98)
P-value, test of diff.	0.000	0.685	0.58	0.329	0.385	0.222	0.234
CCM, Minority		0.692	0.729	0.839	3644	6082	9728
CCM, White		0.753	0.715	0.865	4406	6611	11011

Notes: N = 37,653 Minority youth and N = 5,366 White youth. 390 observations are dropped due to missing race/ethnicity. Earnings winsorization recodes each quarter's highest earnings to the 99.5th percentile of all quarterly earnings before summing across years. CM shows control means; CCM shows control complier means. Regressions include baseline covariates. Standard errors clustered on individual are shown in parentheses. * $p<0.1$, ** $p<0.05$, *** $p<0.01$

Table 8: IV Effects on Joint Employment and Graduation Outcomes for Minority and White Youth

		Panel A: On-Time Graduation			
		Ever Work, On-time	Never Work, On-time	Ever Work, Not On-time	Never Work, Not On-time
	First Stage				
IV, Minority	0.4469***	0.0053	-0.0266**	0.0345***	-0.0130*
	(0.0063)	(0.0151)	(0.0119)	(0.0109)	(0.0075)
IV, White	0.3643***	-0.021	0.0422	-0.0388	0.0259
	(0.0201)	(0.0619)	(0.0563)	(0.0335)	(0.0251)
P-value, test of diff.	0.00	0.681	0.233	0.038	0.139
CCM, Minority		0.721	0.122	0.115	0.041
CCM, White		0.791	0.097	0.115	0
		Panel B: Any Graduation			
		Ever Work, Graduated	Never Work, Graduated	Ever Work, Not Graduated	Never Work, Not Graduated
IV, Minority		0.0276*	-0.0322***	0.0109	-0.0063
		(0.0150)	(0.0122)	(0.0097)	(0.0068)
IV, White		-0.0258	0.0492	-0.0477*	0.0175
		(0.0615)	(0.0570)	(0.0290)	(0.0234)
P-value, test of diff.		0.399	0.163	0.056	0.331
CCM, Minority		0.753	0.131	0.084	0.032
CCM, White		0.811	0.09	0.11	0
		Panel C: Any Graduation or Continued Attendance			
		Ever Work, Persisted	Never Work, Persisted	Ever Work, Not Persisted	Never Work, Not Persisted
IV, Minority		0.019	-0.0251**	0.0199**	-0.0131**
		(0.0150)	(0.0126)	(0.0090)	(0.0062)
IV, White		-0.0307	0.0436	-0.0417	0.0221
		(0.0629)	(0.0583)	(0.0286)	(0.0213)
P-value, test of diff.		0.443	0.25	0.040	0.113
CCM, Minority		0.788	0.132	0.049	0.030
CCM, White		0.849	0.1	0.07	0

Notes: N = 12,589 Minority youth and N = 1,085 White youth. 58 observations in graduation data are dropped due to missing race/ethnicity. Sample and outcomes defined in Table 6. CCM shows control complier means, rounded to 0 if estimate is negative. Regressions include baseline covariates. Standard errors clustered on individual are shown in parentheses. * p<0.1, ** p<0.05, *** p<0.01

Figure 1: Experimental Flow Chart

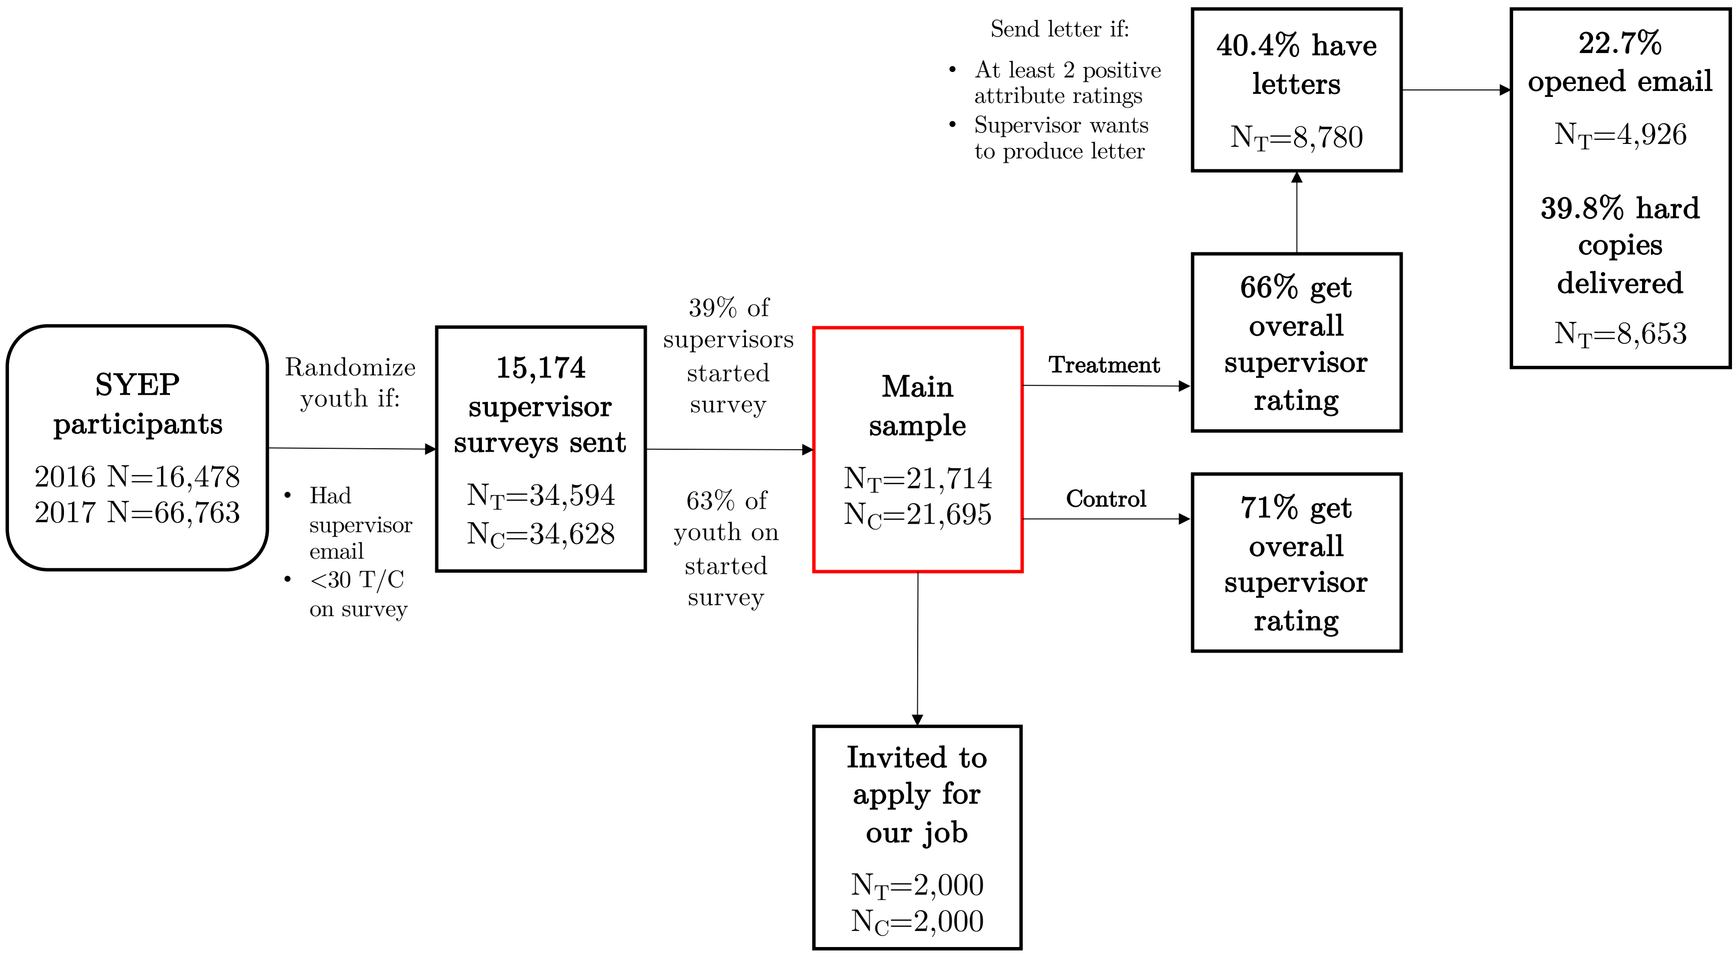

Figure 2: Screenshots about Treatment Youth on Supervisor Survey

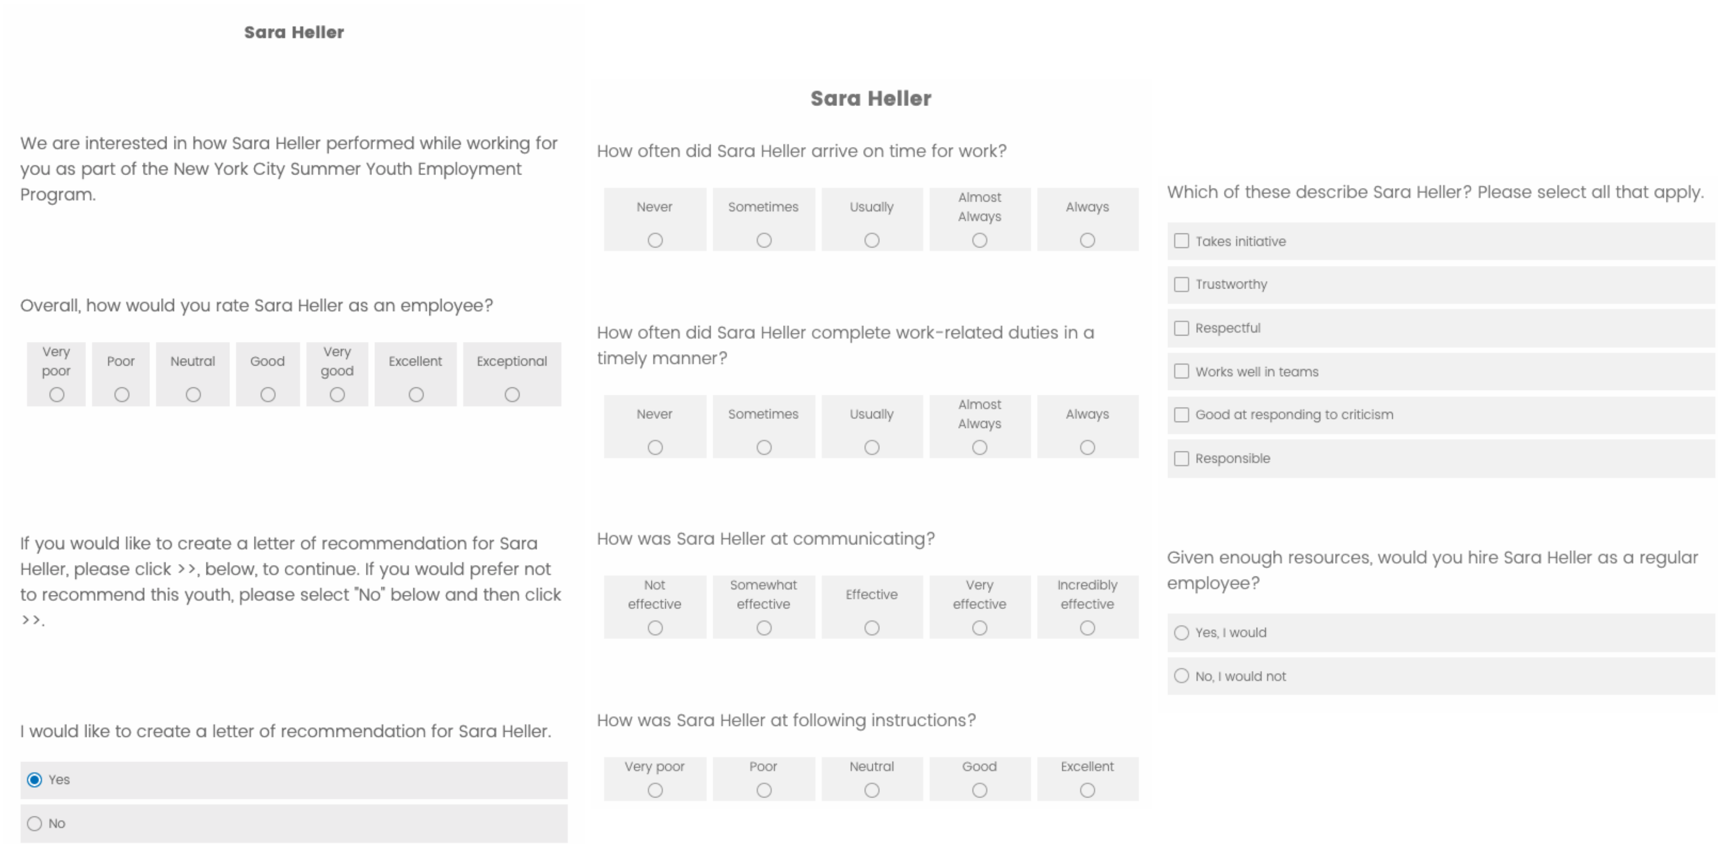
Sara Heller

We are interested in how Sara Heller performed while working for you as part of the New York City Summer Youth Employment Program.

Overall, how would you rate Sara Heller as an employee?

Very poor | Poor | Neutral | Good | Very good | Excellent | Exceptional

If you would like to create a letter of recommendation for Sara Heller, please click >>, below, to continue. If you would prefer not to recommend this youth, please select "No" below and then click >>.

I would like to create a letter of recommendation for Sara Heller.

- Yes
- No

Sara Heller

How often did Sara Heller arrive on time for work?

Never | Sometimes | Usually | Almost Always | Always

How often did Sara Heller complete work-related duties in a timely manner?

Never | Sometimes | Usually | Almost Always | Always

How was Sara Heller at communicating?

Not effective | Somewhat effective | Effective | Very effective | Incredibly effective

How was Sara Heller at following instructions?

Very poor | Poor | Neutral | Good | Excellent

Which of these describe Sara Heller? Please select all that apply.

- Takes initiative
- Trustworthy
- Respectful
- Works well in teams
- Good at responding to criticism
- Responsible

Given enough resources, would you hire Sara Heller as a regular employee?

- Yes, I would
- No, I would not

Notes: The image on the left shows the first screen supervisors saw asking about each youth with the overall rating question and the invitation to write a letter. As indicated in the image, the option to create a recommendation was pre-selected. The images in the middle and on the right show the questions asked about each treatment youth when the supervisor agreed to create a letter of recommendation.

Figure 3: Example Letter of Recommendation

Department of
Youth & Community
Development

November 1, 2017

To Whom It May Concern:

Sara Heller worked for me at the Wharton School during the summer of 2017. Overall, Sara was an exceptional employee.

With regard to reliability, Sara was always on time to work. Sara always completed work related tasks in a timely manner.

When it came to interpersonal interaction, Sara was an incredibly effective communicator. Sara was excellent at following instructions.

In addition to Sara's other strengths, Sara takes initiative, is trustworthy, is respectful, works well in teams, is good at responding to constructive criticism, and is responsible.

Given the resources, I would hire Sara as a full-time employee. I invite you to contact me if you would like more information. I can be reached at 215-898-7696 or judd.kessler@wharton.upenn.edu.

Sincerely,

Judd Kessler
The Wharton School

The New York City Department of Youth and Community Development (DYCD) invests in a network of community-based organizations and programs to alleviate the effects of poverty and to provide opportunities for New Yorkers and communities to flourish.

Empowering Individuals • Strengthening Families • Investing in Communities

Note: This recommendation letter is part of a pilot program being run by the New York City Department of Youth and Community Development. Some youth were randomly selected to be part of the pilot. These youth were eligible to receive a letter of recommendation, which reflects supervisor feedback about each individual's job performance.

Notes: See Figure 2 for the source of inputs into each sentence for this example letter.