LEARNING TO QUIT? A MULTI-YEAR FIELD EXPERIMENT WITH INNOVATION DRIVEN ENTREPRENEURS*

Esther E. Bailey
Daniel C. Fehder
Eric J. Floyd
Yael V. Hochberg
Daniel J. Lee

June 10, 2024

We use a randomized experiment with 553 science- and technology-based startups in 12 co-working spaces across the US to evaluate the effects of intensive, short-term entrepreneurial training programs similar to that offered by accelerator programs and executive education programs in US business schools on survival and performance for innovation-driven startups. Treated startups are more likely to shut down their businesses and do so sooner than controls. Conditional on survival, however, treated startups are more likely to raise external funding for their ventures, raise funding faster, and raise more funding than the control group; they also exhibit higher employment. Treated founders are less likely to found a new startup after shutdown. Our findings are consistent with practitioner arguments that early entrepreneurship training interventions can help entrepreneurs with less viable ventures "optimally quit" ("fail fast"). We use machine learning techniques (causal random forest) to provide initial insights on the most impacted subgroups.

Keywords: Entrepreneurship, Fast Failure, Optimal Quitting, Field Experiments, Education

JEL Classification: C93, D22, M113, M53, O32

---

* We thank conference and seminar participants at Bocconi University, Duke University, Georgia Tech, Harvard Business School, McGill University, Michigan State University, Purdue University, University of Chicago, University of Colorado, University of Mannheim, University of Southern California, University of Utah, University of Washington, the Advances in Field Experiments Conference, and the Southern California Strategy Conference for insightful feedback on both experiment design and/or findings. The Kauffman Foundation provided generous financial support for this project under Foundation Grant G-201508-20151274. All errors and omissions are our own. Bailey: University of Houston; Fehder: University of Southern California; Floyd: University of California at San Diego; Hochberg: Rice University & NBER; Lee: University of Delaware. A randomized controlled trials registry entry is available at AEARCTR-0001566 (Fehder et al., 2016). This study was conducted under the oversight of the Rice University Institutional Review Board (IRB).

1. INTRODUCTION

Innovation-driven entrepreneurial activity is considered a major driver of economic growth. Venture capital-backed startups, which typically fall into this category, account for 41% of total US stock market capitalization and 62% of publicly traded companies' research and development spending (Gornall and Strebulaev, 2021). Policies that enable the entry of new innovation-driven entrepreneurs in a region can spur future growth in both venture capital allocation and employment, potentially yielding a virtuous cycle of growth (Hausman 2022; Hausman, Fehder, and Hochberg 2023). As a result, policy makers often undertake significant efforts to encourage and support entrepreneurial activity. At the same time, most new innovation-driven ventures fail, in part because many entrepreneurs may only have experienced the technical side of the labor force. Importantly, many innovation-driven entrepreneurs may not have sufficient knowledge to fully understand the entrepreneurial process, challenges they may face, or how to evaluate the likelihood of success. While post-entry learning and subsequent exit are central to many economic models of entrepreneurship (Jovanovic 1982; Vereshchagina and Hopenhayn 2009), the efficiency of the learning process and its connection startup performance and exit are less explored (Manso 2016). Increasingly, the stated goals of entrepreneurship training and acceleration programs are twofold: not only to improve and accelerate success for viable ventures, but also to encourage "fast failure" for non-viable ventures or entrepreneurs (Cohen et al. 2019).[1] Whether such programs achieve either of these goals, however, remains an open question.

In this paper, we assess the potential for such short-term entrepreneurship training to affect innovation-driven entrepreneurship (IDE) startup "rational quitting" and venture success. From a theoretical perspective, the impacts of entrepreneurship training are ambiguous. A natural null hypothesis is that training should make entrepreneurs "better,"

[1] An example of a startup accelerator program is Techstars, whose graduates have raised $25.7B in total funding and created $98.6B in market capitalization. These programs offer short, intensive educational programs to inform founders about the entrepreneurial process, help evaluate ideas and markets, and aid in raising financing. See Hochberg (2016) for a detailed definition and description of such programs.

which is often interpreted as prolonging startup survival. On the other hand, treated startups could shut down faster if training helps entrepreneurs realize their ventures do not warrant additional investment of time or resources, either because of weaknesses in the venture, product, team, or market, or because the founding team did not wish to continue pursuing an entrepreneurial venture once they had more knowledge of the process and challenges in hand—a form of "rational quitting."

From a practical standpoint, identifying a causal impact of training interventions is also challenging. To address causality, random variation is needed, yet there are few settings in which entrepreneurship training is randomly or exogenously provided. For example, MBA programs, which often house entrepreneurship curricula, are not filled with randomly selected students, nor are entrepreneurship-specific certificate or executive education programs. Quasi-experiments resulting from random shocks to training are similarly uncommon. Here, instead, we implement a multi-year field experiment with a large set of operating IDE startups, who were screened ex-ante to have high growth potential. As part of the study, we randomized a training intervention that provides initial coverage of the frameworks contained in an MBA curriculum in entrepreneurship condensed into six four-hour overview sessions conducted over consecutive weeks. The total course time is comparable to the instruction time in a nine-week (i.e., quarter-long) class at business schools.

The short term, intensive nature of the curriculum was not chosen ad hoc. Rather, it was identified ex-ante as the most intensive curriculum that could be offered to active entrepreneurs operating tech- and science-based going concerns that would conceivably allow them to find time to participate, and which could be scaled post-experiment, if successful. The training program mimics the overview nature of entrepreneurial training sessions offered as part of many startup accelerator programs and executive education short courses. The goal of the condensed curriculum was to introduce participants to key entrepreneurship-specific knowledge.

Our randomized control trial (RCT) spanned metro areas across the US, partnering with 12 startup incubators and co-working facilities that provide workspace to technologyand science-based startups, and was conducted over a period of four years. We marketed our efforts as the “Entrepreneurial Success Initiative” (ESI). Participating startups were randomized within each participating site. Treated startups received the training program and a $1,000 participation incentive paid to their corporate account. Control group startups received only the participation incentive. We collected extensive information on the startups and their management teams at baseline and conducted periodic surveys at six-month intervals thereafter.[2]

We examine the effect of the entrepreneurship training program across six categories of interim outcomes: survival; pivoting to a new customer segment, business model, or product offering; external fundraising; maximum employment the startup reached; and individual founder outcomes post-startup shutdown (employment vs entrepreneurship). How does training affect startup outcomes? Our experiment reveals a number of intriguing findings. First, treated entrepreneurs are *more* likely to shut down their ventures and do so *sooner* than the control group startups. Of course, a startup may, rather than choosing to shut down, instead choose to change either their market, customer, business model, or product offerings, a practice often referred to as “pivoting” (Ries 2011). When we estimate competing hazard models for pivoting versus shutdown, the coefficients are directionally consistent with lower hazard of pivoting ($p=0.09$ to 0.14), suggesting that treated startups may be choosing to shut down rather than experiment with alternatives when faced with difficulties.

There are a number of ways to interpret these findings. One is as a negative effect: training could harm the startup or entrepreneur and lead to faster shutdown by either discouraging the entrepreneur, wasting their time, or confusing them in some manner. Alternatively, the training may help entrepreneurs better recognize their venture’s prospects and the skills and behavioral characteristics required for success, and lead those with poorer prospects to shut down sooner.

[2] Individual founders and management team members participated in the data collection on a voluntary basis under IRB supervision and were free to withdraw from the experiment at any time.

To distinguish between the two, we examine fundraising and employment. The ability to raise funding from outside investors, typically venture capitalists, is often used in the financial economics literature as a signal of interim success and quality for IDE ventures. Overall, the fundraising empirical patterns support the notion that our results reflect "rational quitting"—what the startup ecosystem often refers to as "fast failure"—rather than harm. Both unconditionally, and accounting for the competing outcome of startup shutdown via Fine-Gray competing risk models and Cragg double-hurdle estimators, we find that treated startups are more likely to raise external capital from venture capitalists, raise capital sooner, and raise more funding dollars than do control group startups. Consistent with our funding results, when we measure employment differences between treated and control startups, we find that treated startups reach maximum employment levels that are between 21-28% higher than control startups.[3]

Overall, the findings suggest that treated firms that do not shut down are, on average, of *higher* quality than surviving control group firms; even accounting for the early shutdown effect, treatment improves the treated startups' ability to raise funds and grow in employment. Importantly, the empirical patterns do not appear to be consistent with training merely discouraging entrepreneurs overall; if the mechanism for shutdown is discouragement, it appears to be of startups with poorer prospects, on average.

We then turn to post-shutdown labor market choices. If training helps entrepreneurs better understand how to assess the viability of a startup or better evaluate their suitability for entrepreneurial activity, we might expect treated entrepreneurs to change their subsequent labor market choices: they may raise their bar for entry into entrepreneurship in the future or change their evaluation of their expected value from a startup relative to their outside option. Utilizing information on the post-shutdown labor market activities of founders, we show evidence consistent with this hypothesis:

[3] We focus on maximum employment because innovation-driven startups frequently grow quickly but then shutdown when market and technical information revealed by operating changes their evaluation of startup's future value relative to their outside option (Kerr, Nanda, and Rhodes-Kropf 2014).

conditional on having left their current startup, founders of treated startups are less likely to immediately pursue another startup as their next job.

Since attendance at the training sessions was not mandatory, our estimates represent "intent to treat" (ITT) rather than "treatment on the treated" (ToT). We utilize attendance data to estimate the TOT. In the first stage, we instrument two measures of attendance (attended at least one session and number of sessions attended) using the randomly assigned treatment. We then utilize those estimates to calculate the ToT effect for survival, fundraising, maximum employment, and post-shutdown labor market choices. As expected, the magnitude of the TOT (attended at least one session) is somewhat larger than the ITT magnitude.

Finally, the overall treatment effect may mask important heterogeneity. Using theory to predict the effects of moderating variables on the treatment in our setting is subject to (at least) three limitations. First, there is an absence of theory that explicitly identifies which variables we would expect to moderate the effects. Second, heterogeneity is not randomly assigned in the same way as treatment status. As a result, it is difficult to use heterogeneous effects to make causal inferences. Third, if the researcher only considers variables to include in this type of analysis ex-ante, they may unintentionally ignore important variation in the data ex post.

To this end, we implement a causal random forest approach to detect heterogeneous treatment effects (Wager and Athey 2018; Athey, Tibshirani, and Wager 2019). The algorithm allows us to examine a large set of potential subgroups and identify patterns in the data that are potentially hidden by the local average treatment effect (LATE). Our analysis suggests heterogeneity for the treatment effect on startup survival, fundraising, and employment. For survival, only an indicator for whether a founder has a top-10-degree moderates treatment. In contrast, a larger set of characteristics appears to moderate the effect of treatment on fundraising, including incorporation status, prior startup experience, and prior P&L experience. For employment, startups with a STEM educated founder moderates the treatment effect. Importantly, this analysis is descriptive, not causal. Future research will be necessary to address potential causal implications.

Our paper contributes to several literatures. First, we contribute to the interdisciplinary topic of "rational quitting." Although many economic models suggest that rational economic agents should quit when their outside options exceed the benefits of their project (that is, the NPV of their project is negative), few studies have been able to empirically show interventions that can increase the propensity of agents to behave in line with economic theory. "Rational quitting" serves as an important counterpoint to many mainstream dialogues about the unequivocal benefits of perseverance (List 2022). A growing literature has documented how entrepreneurs face substantial information frictions when trying to assess the quality of their business ideas (Lerner and Malmendier 2013; Scott, Shu, and Lubynsky 2019; Bennett and Chatterji 2023). Our study suggests the possible presence of broader information frictions that may impact an entrepreneur's understanding of the growth process, resulting in reduced ability to identify when it is optimal for them to quit.

Second, our paper contributes to the literature on the efficacy and effects of targeted interventions for entrepreneurs. The economics literature in this area has primarily focused on interventions in the developing world (e.g. McKenzie 2017; De Mel, McKenzie and Woodruff 2019; Davies et al. 2024), often involving either SMEs (e.g. Bruhn, Karlin, and Schoar 2018) or subsistence entrepreneurs (e.g. Karlin and Valdiva 2011). Recent studies in the management literature span both the developing and developed world, and include the effects of posting non-pecuniary versus pecuniary messaging when recruiting participants to an entrepreneurship program (Sen, Guzman, and Oh, 2018), the effect of communications and networking assistance for small business performance (Kotha el al., 2023), the impact of training on methods of scientific experimentation on entrepreneurial entry and revenues conditional on entry for would-be entrepreneurs across a wide variety of business types (Camuffo et al, 2019), and the effects of availability of peer networking for startup survival and growth (Chatterji et al, 2019).

Our study places specific focus on IDE entrepreneurs running active technology-related businesses with significant growth aspirations—a group of significant interest to policy makers and private funders that spend substantial resources on education for entrepreneurs each year [4] —and explores the effect of broad coverage IDE-entrepreneurship-specific training on outcomes over a substantial time period (two years). Our study provides unique insights on the viability of general entrepreneurship training programs in developed innovation ecosystems characterized by high human capital and financial capital. If entrepreneurship training can reduce inefficient continuation and boost the performance of higher-potential startups, policymakers can potentially harness training interventions to encourage growth and business development. On the other hand, if educating entrepreneurs is mostly ineffective, knowing this could save considerable resources by deterring policymakers from focusing on an intervention that serves little purpose. Our findings suggest that education for particular sets of entrepreneurs may help them better assess the potential for their ventures and prevent inefficient continuation.

Finally, a sizeable proportion of field experiments in economics and business-related fields have been implemented in developing economies. For IDE entrepreneurship in developed countries, it is unclear that results would generalize from studies in developing countries. Notably, one common challenge to conducting entrepreneurship field experiments in developed countries is finding settings that allow for appropriate population sample sizes with suitable power to detect effects. Here, we achieve this through partnership with practitioner organizations. Our RCT demonstrates that partnerships and ongoing relationships with practitioners can be an important solution for achieving scale in these settings. Addressing our research equation of interest would not have been possible within a single tech-oriented co-working facility alone and required the participation and partnership of eleven different organizations (twelve different co-working spaces).

2. BACKGROUND AND EXPERIMENTAL DESIGN

Our analyses represent the first set of results from the RCT registered under AEARCTR-0001566 (Fehder et al., 2016). We recruited innovation-driven startups from

[4] McKenzie (2021) estimates that more than $1 billion is spent on entrepreneurship training each year. For example, the US government spends roughly $40 million on iCorps training yearly for IDE entrepreneurs.

12 startup-focused co-working entrepreneurial spaces in ten metro areas (eleven cities) across the US. Importantly, participating locations did not include cohort-based accelerator programs or other programs that offer financing or structured programming for startups.[5] Rather, we used startup-oriented co-working spaces that at most offered ad hoc workshops.

2.1. Participant Population

Interested startups were screened to confirm that: (i) the startup was IDE (this excludes businesses such as food trucks, marketing agencies, development shops, consultancies, and other non-tech or non-science-based businesses); (ii) they had not yet raised any institutional funding (firms that had previously raised venture capital or other institutional financing were restricted from participating); and (iii) their founders and top management team members (when not founders) were willing to participate in the field study data collection effort, which included the randomized entrepreneurial training program offered as a participation incentive under the terms of the IRB-approved study.[6]

At the corporate entity level, participating startups signed contracts which paid $1,000 up front to the company legal entity (not the founders) in exchange for an obligation on behalf of the corporate entity to fill out a series of experimental surveys administered by the research team over a multi-year period. The baseline survey for the startup firms includes background information on the company and its activities and performance to date. At the individual level, members of the management teams of the contracted startup companies participated in the research experiment under IRB supervision, with the option for the individuals to withdraw from the study at any time. Baseline information collected includes demographic, education, and other information such as behavioral characteristics.

Follow-up surveys sent to the participants every six months included questions about the startup and its ongoing performance. The company survey questions tracked standard

[5] For a discussion of the differences between co-working spaces, incubators and accelerators, see Cohen et al. (2019).

[6] Willingness to participate did not require individual members of the team to show up for training if their startup was randomized into treatment but did require them to fill out the baseline individual survey.

measures of interim startup success such as funding, employment, team composition, and various corporate practices. Both the baseline and follow-up surveys were designed to be as comprehensive as possible while also deferring to the realities of participant time constraints and burden of completion.

Figure 1 presents a map of treatment sites, and Table 1 presents further details on participating incubators and co-working spaces. The cities covered in the RCT were deliberately chosen to be a mix of more established startup ecosystems (Boston, Austin, New York City) and newer startup ecosystems (e.g. Chicago, St. Louis, Houston), while co-working spaces were chosen with an eye towards significant tech- and science-based startup membership, but deliberately avoided locations that themselves offered formal training sequences for their lessees.

A between-survey period of six months was chosen so that we could closely track startups as they evolved while simultaneously allowing enough to time to pass for reasonable changes to occur. Survey response rates over time are presented in Table 2. Participation rates for the surveys were high, notably so compared to most prior survey work. For example, at the first follow-up survey, our lowest percentage of respondents was 82.1% (Los Angeles) while our highest was 100% (Saint Louis and Houston). We did not run the experiment in the San Francisco Bay Area due to the extreme prevalence of local entrepreneurship related programming, support organizations, and IDE social networks. To achieve this, we made use of the claw-back provision in the spirit of Fryer et al. (2022), which allowed us to demand reimbursement for up to the full $1,000 if startups did not fulfill their contractual survey obligations. Survey participation rates in the full sample fall over time, as startups attrite from the sample due to failure. Conditional on survival, response rates at the fourth follow-up (24 months) still reached approximately 80%. In Online Appendix Table A6, we show the results of analysis of differential attrition rates and find no statistically significant difference between the response rates for follow-up surveys except for the first.

The quality of responses to the survey questions, conditional on filling out the survey, exhibited more variability, which was expected. Certain types of written answers required cleaning to ensure that the answers recorded are logical and consistent. We describe the full timing of the experiment across sites and data processing and cleaning procedures in Online Appendix A. The resulting dataset is a panel of 553 firms over four calendar years.

2.2. Intervention

We randomized a condensed version of a standard innovation-driven entrepreneurship (IDE) curriculum to the startups in our sample. This consisted of six sessions of four hours each, delivered at each site over six consecutive weeks. An ex-ante concern is the large set of potential interventions; if we did not find a meaningful treatment effect with a particular intervention, it would not preclude the possibility that other potential treatments could produce a treatment effect. Furthermore, even if we did find a treatment effect, other interventions could produce even more pronounced effects. Given these considerations, our intervention was designed based on feedback from a pilot study. The curriculum was chosen to be long enough to convey a useful amount of knowledge (particularly critical information the inexperienced might not be aware of) but condensed enough that startup teams would participate given extensive demands on their time. Importantly, the modules covered were a condensed overview of typical entrepreneurship content taught in top business school executive education programs and other training programs (such as accelerator programs). A more detailed description of the curriculum is provided in Appendix X.[7]

2.3. Randomization

We randomized startups into treatment and control *within* each site rather than placing some co-working spaces into treatment and some into control. Consequently, each site had startups that were part of both the treatment and control. The disadvantage to this approach is the possibility of spillover from the treatment to control (i.e., the treated

[7] The study was not powered or designed to assess to detect the differential impact of the different subject matter addressed in the curriculum; treated entrepreneurs were encouraged to attend all sessions and overwhelmingly did so.

group could, in theory, teach the control group what they had learned), which would bias against us finding differences. The advantages include increased statistical power and better control for the specific characteristics of any given city or co-working space, which outweighed this disadvantage. Due to capacity constraints on classrooms for training, in most locations, fewer startups were randomized into treatment than into control.

Randomization occurred after startups enrolled in the program, in order to avoid treatment specific selection bias (Floyd and List, 2016). Startups in the treatment group, however, were not contractually obligated to send their management team members to the training sessions; individual participation in the experiment was voluntary. As a result, our average effects represent intent-to-treat. We supplement our analysis with attendance data to better identify the potential effect of the treatment on the treated. A summary schematic of our field experiment design is presented in Figure 2.

3. DATA

The experiment resulted in an extensive set of data. Given the scope of our analyses in this paper, we focus on only a subset. We utilize the following outcome variables: (1) **Time to closure for a startup:** defined as the total number of days from the date in which the startup enrolled in our study to the date in which the startup was no longer operating. A startup is coded as shut down either based on survey response or given an external signal of closure such as the first point in time that every (former) employee of the firm on LinkedIn has a job somewhere other than the startup. We explain the construction of our time to closure variable in Online Appendix A3; (2) **Pivots:** major changes to either the customer, product, or revenue/distribution channel. Details of how we determine whether startups pivot are in Online Appendix A3; (3) **Funding amount raised:** the total dollar amount of funding that a startup raised, obtained from Pitchbook in order to ensure a standardized measure for all participants without attrition concerns; (4) **Employment:** the maximum number of employees working at the startup during the survey period; (5) **Post-shutdown founder labor choices:** whether founders launched another startup immediately after leaving the ESI enrolled startup.

We include a number of additional variables. These include basic behavioral and demographic characteristics such as gender, risk preferences, and incorporation status. This data is taken directly from our surveys. Variable names and definitions are presented in Table 3. Table 4 presents the overall summary statistics for the treatment and control sample at baseline, and Table 5 presents summary statistics separated by treatment and control along with a test of balance. The first column of Table 5 presents the mean and standard deviation (in parentheses) of the baseline covariates of the startups in the control group. The second column shows the mean and standard deviation (in parentheses) of the baseline covariates of the startups in the treatment group. The third column shows the mean and standard deviation (in parentheses) of the baseline covariates of all the startups in the sample. To test for sample balance, we estimate separate regressions for each covariate of the form $Treat = \beta_0 + \beta_1\ Covar + \epsilon$. In the column, we show the estimates of β_1 for each covariate. Randomization appears to have worked reasonably well given our sample size. Of the sixteen baseline variables presented, only two, Corporation and Business Education, demonstrates statistically significant differences between the treatment and control groups. The existence of some differences is to be expected given the dimensionality. To account for any residual differences post-randomization, we control for a variety of baseline variables in our regressions.

We next turn to individual founder behavioral parameters. Because we measure behavioral parameters at the founder/individual team member level, we must first determine how to aggregate parameters to arrive at a single value for the startup. For simplicity, we use the maximum of each behavioral characteristic across the startup management team.[8] Risk preferences are measured on a scale of 1 to 4, following Barsky et al. (1997), with a 4 being most risk loving. Our firms have a mean of 3.60 compared to a mean of 1.72 in Barsky et al. (1997), indicating that our sample of entrepreneurs is relatively more risk loving than average, consistent with prior research. Similarly, the max optimism score of almost 20 indicates that our sample exhibits high optimism relative to

[8] Online Appendix B shows robustness to alternative construction of the variables that require the aggregation of individual-level founder variables into startup-level variables.

the mean of 14.33 reported for the standard laboratory population used to develop our measure of optimism (LOT-R, Scheier, Carver, and Bridges 1994).

4. MAIN RESULTS

Because our startups were only contracted to a two-year follow-up reporting period, and due to the often-extended time period it takes to exit an IDE startup, we necessarily cannot follow all firms through to their end outcome. Instead, we focus on interim performance measures indicative of startup progress and performance. For some variables, we can supplement data collection beyond the experimental survey instruments, as described in Section 3, allowing us to extend our analysis period outside of the two-year follow-up period.

4.1. Survival

A natural null hypothesis is that training should make entrepreneurs "better," thus prolonging their startups' survival. On the other hand, treated startups could shut down faster because training encourages a form of "rational quitting." In Table 6, we first report the raw percentage of startups still in business at different time periods. Unsurprisingly, a large number of startups shut down over the study period. Overall survival at three years post intake is 30.2%. At the three-year mark, 34% of control firms are still operating, while only 26% of treated firms remain in business. This difference is economically large. In the raw survival data, treatment startups shut down operations approximately 105 days sooner on average than control firms.

Figure 3 Panel A shows Kaplan-Meier plots for survival. It is visually apparent that the difference in survival rates between treatment and control in our sample manifests relatively early—within approximately one year. A Wilcoxon rank-sum test confirms the treatment and control groups do not have the same survival functions ($p<0.0001$). Survival in our analysis is restricted to the narrow window (3 years) in which we observe startup outcomes.

Kaplan-Meier survival curve analysis does not allow for the analysis of survival differences conditional on control variables. We next estimate a series of Cox proportional hazards models for survival (time to shutdown), where the independent variable of interest is treatment. We estimate variations on the following model:

$$h(t|x) = h_0 \exp\left(\beta_1 treated + \gamma_s + \varphi_i + \xi_i + X^{'}\beta\right) \tag{1}$$

where h_0, the baseline hazard rate, is estimated non-parametrically, γ_s is a site-level fixed effect, φ_i is an industry fixed effect, ξ_i is a stage-of-development fixed effect, and X is a vector of covariates that varies by model. The exponentiated coefficient, e^{β_1}, provides an estimate of the impact of treatment with an intuitive interpretation as a hazard rate ratio. To illustrate interpretation, a hazard ratio of 0.5 represents a 50% reduction in the hazard of shutdown, whereas a coefficient of 1.5 would correspond to a 50% increase. Throughout the paper, we show exponentiated coefficients.

Table 7 presents the estimates. Treatment exhibits an encouraging effect on startup shutdown. The exponentiated coefficients (hazard ratio) on the treatment indicator are positive and significant across all models. The estimated magnitudes are fairly stable across specifications and suggest that treatment increases the hazard of shutdown by 57-61%.

4.2. Pivoting

Rather than quit, a startup that feels its current approach is not viable may instead choose to change their business, adjusting either their market, customer, business model, or product offerings to pursue a different path, a practice called pivoting (Reis 2011; Camuffo et al., 2020). Panel B of Figure 3 shows the Kaplan Meier plots of time to first pivot in our sample. In the raw data, control group startups show a higher hazard of pivoting. A log-rank test confirms that these differences are statistically significant (X^2=4.97, p=0.025).

To quantify the impact of treatment more formally, we estimate a series of hazard models measuring differences in the time to first pivot across treatment and control. The estimates are presented in Panel A of Table 8. Treatment reduces the hazard of a first pivot by a statistically significant 33% (column 1), and the coefficient estimate does not change appreciably with the inclusion of company-level covariates (column 2). When we introduce a larger set of controls in columns (3) and (4), the coefficient estimates remain relatively stable and remain statistically significant, although at a lower level.

It is also the case that a startup's choice to shutdown might preclude us from observing them pivoting, creating competing risks. To address this issue, we estimate Fine-Gray models of competing risks. In these models, we consider the risk of our primary event, pivoting, while also considering the competing risk of firm shutdown. The subdistribution hazards of both events are considered while estimating the cumulative incidence function of the primary event using the following model:

$$1 - CIF(t) = \left(1 - CIF_0(t)\right)\exp\left(\beta_1 treated + \gamma_s + \varphi_i + \xi_i + X'\beta\right) \quad (2)$$

where CIF_0, the baseline cumulative incidence function, is estimated non-parametrically, γ_s is a site-level fixed effect, φ_i is an industry fixed effect, ξ_i is a stage-of-development fixed effect, and X is a vector of covariates that varies by model. The exponentiated coefficient, e^{β_1}, provides an intuitive interpretation as a relative change in the subdistribution hazard function for the primary event, —pivoting. Here, an exponentiated coefficient of 1.5 can be interpreted as a 50% increase in the risk of the primary event in subjects that remain event free or experienced a competing event (Austin and Fine 2017).

The estimates from the competing risk models are presented in Panel B of Table 8. Column (1) shows a statistically significant 40% decrease in the risk of pivoting in our treated startups relative to control after accounting for the risk of shutdown. In columns (2)-(4), the coefficient estimates remain relatively similar in magnitude, but no longer maintain standard statistical significance (p-value ranging from 0.107 to 0.137). Given that only 15% of startups in our sample pivot, this stem in part from a lack of power.

The startup survival results suggest substantial differences in how treated and control firms behave. These patterns are consistent with the lessons contained in the training modules. Firms are provided with tools to assess the likely demand for their product or service, as well as provided clear guidance for the growth they will need to secure future funding. Given the content of our training, we believe that it is likely that treated firms reassessed their expected returns to continuing with their current company.

4.3. Fundraising

Our evidence thus far suggests that treatment encourages "fast failure" amongst startups that received training. However, it could also be the case that our education curriculum makes firms worse off by either discouraging them, wasting their time, or confusing them. To assess this, we turn to performance. We focus on fundraising from external investors as a key performance measure because external financing is particularly important for IDE entrepreneurs. Because IDE entrepreneurs are attempting to bring novel products and services into the economy, they require investment of financial capital in advance of receiving significant revenue or even entering the market (Botelho, Fehder, and Hochberg 2023). Here, we focus on three measures: time to first venture capital (VC) round, whether the startup managed to raise any VC funding, and the logged total amount of VC funding received post-enrollment.

Figure 3 Panel C shows the Kaplan-Meier plots of time to first VC round for treatment and control startups. Treated startups acquire their first VC round more quickly and in higher proportion. A log-rank test of the difference in the survival curves between the two groups is statistically significant ($\chi^2 = 10.41$, $p < 0.001$).

Next, we analyze the impact of treatment on the hazard of a venture capital round using a Cox Proportional Hazard model as in equation (1) and then incorporate the competing risk of shutdown using the Fine-Gray Competing Risk model as in equation (2). The estimates are presented in Panel A of Table 9. Columns (1)-(4) present estimates from the Cox Proportional Hazard models. We find that treatment is associated with a statistically significant 98%-113% increase in the hazard of a startup receiving funding across the different models. Columns (5)-(8) present estimates from the Fine-Gray Competing Risk models. Here, the estimates suggest a statistically significant 69%-81% increase in the risk of receiving venture capital after accounting for the competing risk of startup shutdown. Overall, the estimates indicate that treatment significantly impacted the speed at which treated startups were able to receive funding.

We next turn to exploring startup fundraising in more detail. First, we estimate simple unconditional models of the following form:

$$VC_i = \beta_0 + \beta_1 treated_i + \gamma_s + \varphi_i + \xi_i + X'\beta + \epsilon_{is} \qquad (3)$$

where VC_i in this regression is an indicator variable set to one if the startup received any funding by endline and a zero otherwise ($I\{Raise\}$). γ_s, φ_i, and ξ_i are site, industry, and stage-of-development fixed effects and $X^{'}\beta$ is a vector of additional covariates that vary by model. The standard errors in our model are clustered by site. We estimate a linear probability model (LPM) which provides a simple interpretation of β_1 as a percentage point change in the probability of receiving funding. We show robustness of our results to alternative logit specifications in Table B1 of Online Appendix B.

The estimates are provided in Columns (1)-(4) of Panel B of Table 9. The estimates of the impact of treatment are a statistically significant 8.5-10 percentage point increase in the likelihood that a treated startup ever raises funding relative to the controls. Next, we analyze the amount of funding raised. To do so, we use the natural logarithm of one plus VC dollar raised post enrollment as the dependent variable ($Logged\ Funding$). The results of our analysis are provided in Columns (5)-(8) of Panel B of Table 9. Across all the models, the estimates suggest a positive and statistically significant impact of treatment on the total amount of post-enrollment funding. Depending upon the specification, the estimates suggest that treated startups raise between 228-310% more funding than control startups. Our results are robust to an alternative transformation using the inverse hyperbolic sine (Table B2 of Online Appendix B). If training equally increases the chance that startups close earlier irrespective of quality, we would expect to see decreases in funding rates and total funding in treated startups. This is separate from the human capital channel for our educational program which might increase the quality of the surviving startups and thus both the likelihood of raising and the total amount of funding raised.

While the results above show large differences between treatment and control, they do not account fully for several features of the data generating process. First, they do not account for the extensive and intensive margins for raising funds. Second, they do not account for the relationship between survival and fundraising. To address these issues, we jointly estimate the extensive and intensive margin of treatment while incorporating survival time in a double hurdle model (Cragg 1971). The double hurdle model jointly estimates the probability of raising zero funding and the total amount of funding conditional on being a non-zero observation. We follow Cragg (1971) because the model provides a more flexible functional form for estimating the extensive margin relative to other truncated or censored regression models (Tobin 1958; Heckman 1976; Powell 1986).

The estimates are presented in Panel C of Table 9. Across all specifications, treatment has a significant positive effect on the likelihood of raising funds on the extensive margin after also accounting for survival time. As expected, survival time influences both the extensive and intensive margins; that is, the longer a firm survives, the more likely they are to raise funds, and the more funding they raise. Across most specifications, treatment exhibits a positive and statistically significant impact on both the extensive and intensive margins. Coefficient estimates for the effect of treatment on funding raised are somewhat larger than in the unconditional models.

Figure 4 presents a kernel density plot of the distribution of funding for treatment and control. A Kolmogorov-Smirnov test of the equality of the distributions rejects the null of the same distribution of funding ($p=0.061$). The figure qualitatively demonstrates that the control group has a longer left tail of smaller fundraising outcomes whereas the treatment group has a slightly longer right tail. Thus, it appears that startups that raise low levels of funding are "missing" in the treatment group, and that the treatment distribution is slightly shifted to the right relative to the control distribution.

4.4. Employment

To further reinforce our conclusions, we turn next to an additional performance metric, employment. Startups increase employment levels either in response to current revenue growth or in expectation of future revenue growth. Our employment analysis uses the maximum number of employees the startup reaches across the study period, even if it shuts down within the sample period. This allows us to capture growth in the size of the company even if it eventually shrinks on the way to shutdown. Because employment levels across startups and time periods are both close to zero (interquartile range of 3-10) and whole numbers, we use a Poisson model to account for the count nature of the data.

The result of our analysis is presented in Table 10. All coefficients presented in the table are exponentiated and thus represent incidence rate ratios (IRR) where a coefficient of 1 represents no change and a coefficient of 1.2, for example, would represent a 20% increase in employment. Across all models, the coefficients are significant and show higher maximum employment levels for the treated relative to the controls, with increases range from 21-28% across the different models. These results support the conclusion that treated startups perform better than controls in their ability to grow and pursue their objectives.

4.5. Post-Shutdown Entrepreneurial Choices

Thus far, our results have demonstrated that treatment causes startups to shut down earlier, but, conditional on survival, leaves startups of higher quality. The empirical patterns support the notion that this is a rational effect and is inconsistent with the notion that the training intervention made treated entrepreneurs "worse" or simply discouraged them. Rather, the results suggest that treated firms who do not shut down are of higher, rather than lower quality. Our hypothesis is that treated startups shutdown more quickly because founders have a better understanding of the quality level required to grow and succeed. A natural follow up question is whether there are longer term changes in how these startup founders choose to engage in entrepreneurship after the shutdown of their current startup. In general, founders frequently engage in serial entrepreneurship founding new businesses when one fails. The increased understanding of the hurdles associated with growing an IDE business provided by treatment should spillover into their assessment of future ventures. Thus, training may not only lead founders to "rationally quit," it may also lead them to evaluate potential new startup opportunities more stringently. If this were so, we would expect fewer treated startup founders to jump into entrepreneurship again immediately after shutting down their business.

To investigate this, we collect data on founder employment post-shutdown from LinkedIn. Data is collected after the conclusion of the experiment across all sites, and the cessation of follow-on surveying. We code whether founders of shutdown startups in our experiment start a new startup following the termination of their original experiment startup. We then estimate a series of linear probability models similar in structure to equation (3). Our dependent variable measuring the immediate entrepreneurship choices of founders after shutdown of an ESI enrolled startup.

Table 11 presents the estimates from our analysis. Treated entrepreneurs have a statistically significant reduction of 9.2-11 percentage points in the likelihood of pursuing entrepreneurship again immediately after shutdown of their startup. Our results support the idea that the training not only impacts decisions about the current startup, but also how startup founders weigh the costs and benefits of entrepreneurship more broadly.

4.6. Treatment on Treated

Thus far, all of our analyses measure the effects of intention-to-treat. Attendance in the educational program is not mandatory, and not all teams chose to attend. Startups were randomized into treatment, but ultimately there was no mechanism to force treated individuals to take the class. First, as in all human subject studies, individual participants (founders and management team members) retained the right to cease participation at will. While the possibility existed of incentivizing them to attend sessions via a claw-back mechanism that conditioned on the education program, we decided against this approach. A primary consideration in this design choice is that, if our experiment is to inform policy, understanding voluntary participation rates in the program are likely as important as how effective the training is conditional on startups participating.

While the ITT estimates of our educational intervention are the most policy-relevant, we provide additional TOT estimates for several of our key dependent variables. To do this, we use the randomized assignment to treatment as an instrument to predict various measures of attendance in the ESI classes. We draw our attendance data from RSVPs and ex post measures of attendance. Our two measures of attendance are the total number of sessions attended by the startup and whether they attended at least one session. We then use the instrumented attendance variables to estimate TOT.

The results of our analysis are presented in Table 12. In Panel A, we provide our estimates of the first-stage relationship between treatment assignment and two measures of attendance levels. As expected, there is a large and statistically significant relationship between treatment and attendance. The average first-stage F statistics range between 216.4 and 298.22, well above the thresholds to ensure minimal bias (Stock and Yogo 2005).

In Panel B, we provide the results of the instrumental variable analysis for five key dependent variables: (1) time to shut down, (2) whether the startup raised any funding by experiment endline, (3) logged dollar funding amount, (4) the maximum level of employment the startup reached and (5) whether a startup founder immediately started another company after shutdown. Our estimates of TOT for the program largely conform with the ITT findings. Coefficient estimates for treatment effects increase for all outcomes.

4.7. Treatment Heterogeneity

The exploration of heterogeneity for this study is critical. While our education program is similar to programs at major business schools and entrepreneurship training programs provided by accelerator programs, the program itself is reasonably expensive in terms of both founder time and cost of instructors. We do not expect that all entrepreneurs in our study will benefit equally from the training, but it is not clear which entrepreneurs, as identified by their observable characteristics, are most likely to improve because of treatment. Understanding which characteristics are most important can help policymakers better target the education program towards those that would most benefit.

Without predictions from theory, it is not clear which characteristics are likely to provide treatment heterogeneity. There are enough baseline covariates in our sample data that simply including linear interaction terms into our regressions may provide spurious findings of treatment heterogeneity. In the time that has passed since we initially pre-registered the exploration of treatment heterogeneity in our study, advances in post-processing of RCT data, especially through machine learning (ML) tools, have been considerable. They now provide options for assessing treatment heterogeneity across many covariates in a rigorous manner that does not require a theoretical perspective ex ante.

A number of methods have emerged building upon the tools of machine learning to provide robust estimation of treatment heterogeneity (Chernozhukov et al. 2018; Athey, Tibshirani, and Wager 2019). We focus on the causal random forest because it allows efficient estimation of treatment heterogeneity in settings with smaller sample sizes.[9] The causal random forest approach captures complex, non-linear relationships between covariates, including the treatment variable, and outcomes that might be missed by exploring heterogeneity using OLS regressions with linear interaction terms. A number of existing papers provide useful primers on the underlying theory behind causal random forests (Knittel and Stolper 2019; Davis and Heller 2020). Importantly, we consider the following analysis of heterogeneity to be a descriptive exercise, rather than a test of theory.

Assessment of treatment heterogeneity proceeds after first assessing the degree to which the ML prediction algorithms have adequately captured both the overall average treatment effect (ATE) and the systematic deviations from the average treatment effect for different subgroups, the conditional average treatment effect (CATE). For the sake of brevity, we document these "pre-steps" in Online Appendix C and focus on the outputs of the heterogeneity analysis. In particular, Chernozhukov et al. (2018) propose a test of the model's ability to fit ATE and heterogeneity. The test results provided in Table C1 show that our causal random forests provides good estimates of both the ATE and systematic deviations from it for three out of the five dependent variables in our study (funding, survival, and employment). In addition, Online Appendix Figure C1 indicate that there is significant heterogeneity in the treatment effects for these three dependent variables. The key output of these pre-steps for this method is a validated non-linear model of the CATE which produces a separate estimate of CATE for all unique combination of baseline covariates (see Athey et al., 2019 for further description).

To begin our exploration of the relationship between treatment heterogeneity and baseline covariates, we estimate the best linear projection (BLP) of our baseline covariates onto our estimated CATE distribution. Specifically, we estimate the following regression:

$$\hat{\tau}(X_i) = \beta_0 + X'_i\beta + \epsilon_i \tag{4}$$

where $\hat{\tau}(X_i)$ is the estimated CATE given a vector of covariates, X_i. Because the relationship between X_i and $\hat{\tau}(X_i)$ can be complicated and non-linear, β provides a

[9] Ensemble methods like the causal random forest can be particularly efficient because they can randomly subsample estimation and evaluation data sets individually for each tree (Wager and Athey 2018).

simplified representation of these complicated relationships. Each parameter in the equation estimates how a unit increase in the covariate is associated on average with a change in our estimated CATE.

The results are presented in Table 13. The first column presents the estimates for the model of the relationship between treatment and time to shutdown (survival). The estimates suggest significant relationships between the CATE and Founders with Top 10 Degrees, Founders with Children, and Max Founder Optimism. The estimates for the model of the relationship between treatment and funding are presented in column two. Four baseline variables demonstrate a statistically significant linear relationship with the CATE estimates: Baseline Funding, Incorporation Status, Prior Entrepreneurial Experience, and Prior P&L Experience. The estimates for the relationship between treatment and employment are presented in the third column. Only STEM Education demonstrates a significant relationship with CATE.

We can now assess a related question with clear policy implications: would any given variable be useful for targeting if we had to prioritize treatment to different startups? Because the distribution of CATE is complex and non-linear in the distribution of baseline covariates and equation 4 is a simple linear description, the results of Table 13 do not immediately answer this question. To do so, we estimate the Targeted Operator Characteristic (TOC) curve for the variables that show a significant linear relationship (Yadlowsky et al. 2021). The TOC measures the difference in estimated CATE for the top q-th fraction of startups in the distribution of a covariate, X, from the overall estimated ATE. Specifically, the TOC measures:

$$TOC(q) = E[Y_i(1) - Y_i(0) \mid X_i \geq F^{-1}(1-q)] - E[Y_i(1) - Y_i(0)] \tag{5}$$

where $F(q)$ is the distribution of the covariate X in the sample. The TOC formula has a nice visual representation as the area under a curve representing the excess treatment effect for prioritized treatment units (i.e., higher, or lower values of the covariate).

Figure 5 plots the TOC curves as well as 95% confidence intervals for the covariates that observe significant relationships between our CATE estimates for survival and baseline covariates. While three baseline characteristics emerged in the BLP analysis, only one, Top 10 Degree, shows promise for targeting. Startups with at least one founder with a top-10-degree are significantly less likely to shut down in response to treatment relative to the ATE. Figure 6 presents the same analysis for funding. Three of the four variables show substantial potential as variables in which to target treatment. Companies that are not founded as C-Corporations have a substantially lower excess treatment effect, while founders without prior startup experience and prior P&L experience show substantially larger excess treatment effects relative to the population overall. In contrast, baseline funding, while significant in our best linear projection analysis, does not show promise as a targeting covariate. Figure 7 presents the analysis for employment. Founders without STEM Experience have a higher treatment effect relative to the ATE.

While these methods provide an initial look at potential sources of treatment heterogeneity, we caution that they are descriptive, not causal. Nevertheless, many of the relationships are intuitive and support the motivations behind entrepreneurial training curriculums. For example, substantial excess treatment effects on startup founders with less startup experience or managerial experience suggests that the intervention is affecting the population that most policymakers would ex ante expect it to. Further work, likely experimental, will be needed to better identify and explore treatment heterogeneity.

5. CONCLUSION

Despite its potential importance for stimulating economic growth, little is known about the factors that contribute to innovation-driven startups' success or failure. In this paper, we use a field experiment in combination with a rich dataset collected from surveys of participating startups to investigate the importance of realistic-scope short-course entrepreneurship training on future performance for a sample of innovation-driven startups. Our results suggest that policymakers can use short education curriculums to improve outcomes for startups—a stated goal of many foundations and local development programs. Moreover, our analysis of treatment heterogeneity provides initial information on the characteristics of startups and entrepreneurs that policymakers should prioritize the treatment towards in order to maximize the efficacy of these type of programs.

In addition, our experiment was designed specifically to easily scale towards broad implementation of our program. This design choice is in line with a recent push in the economics literature to consider the potential for scaling of offerings within the context of experimental interventions (List 2022; Al-Ubaydli, Lai, and List 2023). Often, considerable uncertainty remains as to whether experimental interventions can be scaled beyond the scope of the smaller scale randomized study. Our funding source was interested in the effects of education programs that could be potentially offered to a much larger audience.[10] To ensure that our chosen intervention was implementable at scale post-experiment at reasonable cost, we conducted extensive qualitative research at the pilot stage. The final educational program was selected based on its ability to realistically be offered to working entrepreneurs outside of experimental conditions and in the absence of participation incentives. Furthermore, we explicitly considered the ability of the intervention to be scaled through a train-the-trainer model. We utilized corporate trainers who had experience teaching short entrepreneurship training classes who were provided materials and content for the modules that were taught. In other words, we sought to ensure that none of the instructors was sufficiently unique that their main contributions couldn't be replicated by others.

Pinning down the overall welfare consequences of our intervention is difficult. Our results support the interpretation that we cause less viable startups to shut down earlier. Given the earning potential and high prior wages of individuals involved in innovation-driven ventures, it may be beneficial to society if training causes unproductive entrepreneurs to preserve their time and return to the labor force. If successful, certain innovation-driven startups by their nature may lead to outsized societal and economic impact. Such startups are inherently risky, and luck may play a large role in eventual success. We cannot rule out the possibility that some ventures that chose to shutdown sooner as a result of treatment would have ultimately been successful. From a societal perspective, it may instead be optimal to encourage all startups to continue, despite the

[10] In this sense, we were not interested in exploring interventions that might produce a larger treatment effect but were outside the scope of what could reasonably be offered to a broad set of entrepreneurs.

possibility of entrepreneurs wasting time and resources, if policy makers' primary focus is to maximize the probability of increasing innovation. We consider this discussion to be outside the scope of our experiment, and worthy of future research.

REFERENCES

Al-Ubaydli, Omar, Chien-Yu Lai, and John A. List. 2023. "A Simple Rational Expectations Model of the Voltage Effect." Working Paper. Working Paper Series. NBER.

Athey, Susan, Julie Tibshirani, and Stefan Wager. 2019. "Generalized Random Forests" 47 (2): 1148–78.

Austin, Peter C., and Jason P. Fine. 2017. "Practical Recommendations for Reporting Fine-Gray Model Analyses for Competing Risk Data." *Statistics in Medicine* 36 (27): 4391–4400.

Barsky, Robert B., F. Thomas Juster, Miles S. Kimball, and Matthew D. Shapiro. 1997. "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study." *The Quarterly Journal of Economics* 112 (2): 537–79.

Bennett, Victor M., and Aaron K. Chatterji. 2023. "The Entrepreneurial Process: Evidence from a Nationally Representative Survey." *Strategic Management Journal* 44 (1): 86–116.

Botelho, Tristan L., Daniel Fehder, and Yael Hochberg. 2023. "Innovation-Driven Entrepreneurship." National Bureau of Economic Research.

Camuffo, Arnaldo, Alessandro Cordova, Alfonso Gambardella, and Chiara Spina. 2019. "A Scientific Approach to Entrepreneurial Decision Making: Evidence from a Randomized Control Trial." *Management Science* 66 (2): 564–86.

Chernozhukov, Victor, Mert Demirer, Esther Duflo, and Ivan Fernandez-Val. 2018. "Generic Machine Learning Inference on Heterogeneous Treatment Effects in Randomized Experiments, with an Application to Immunization in India." NBER.

Cohen, Susan, Daniel C. Fehder, Yael V. Hochberg, and Fiona Murray. 2019. "The Design of Startup Accelerators." *Research Policy* 48 (7): 1781–97.

Cragg, John G. 1971. "Some Statistical Models for Limited Dependent Variables with Application to the Demand for Durable Goods." *Econometrica* 39 (5): 829–44.

Davis, Jonathan MV, and Sara B. Heller. 2020. "Rethinking the Benefits of Youth Employment Programs: The Heterogeneous Effects of Summer Jobs." *Review of Economics and Statistics* 102 (4): 664–77.

Fryer, Roland G., Steven D. Levitt, John List, and Sally Sadoff. 2022. "Enhancing the Efficacy of Teacher Incentives through Framing: A Field Experiment." *American Economic Journal: Economic Policy* 14 (4): 269–99.

Hausman, Naomi. 2022. "University Innovation and Local Economic Growth." *Review of Economics and Statistics* 104 (4): 718–35.

Hausman, Naomi, Daniel C. Fehder, and Yael Hochberg. 2023. "The Virtuous Cycle of Innovation and Capital Flows," 50.

Heckman, James J. 1976. "A Life-Cycle Model of Earnings, Learning, and Consumption." *Journal of Political Economy* 84 (4, Part 2): S9–44.

Hochberg, Yael. 2016. "Accelerating Entrepreneurs and Ecosystems: The Seed Accelerator Model." *Innovation Policy and the Economy* 16 (1): 25–51.

Huber, Laura Rosendahl, Randolph Sloof, and Mirjam Van Praag. 2014. “The Effect of Early Entrepreneurship Education: Evidence from a Field Experiment.” *European Economic Review* 72: 76–97.

Jovanovic, Boyan. 1982. “Selection and the Evolution of Industry.” *Econometrica* 50 (3): 649–70.

Kerr, William R., Ramana Nanda, and Matthew Rhodes-Kropf. 2014. “Entrepreneurship as Experimentation.” *Journal of Economic Perspectives* 28 (3): 25–48.

Knittel, Christopher R., and Samuel Stolper. 2019. “Using Machine Learning to Target Treatment: The Case of Household Energy Use.” Working Paper Series. NBER.

Kotha, Reddi, Balagopal Vissa, Yimin Lin, and Anne-Valérie Corboz. 2023. “Do Ambitious Entrepreneurs Benefit More from Training?” *Strategic Management Journal* 44 (2): 549–75.

Lerner, Josh, and Ulrike Malmendier. 2013. “With a Little Help from My (Random) Friends: Success and Failure in Post-Business School Entrepreneurship.” *The Review of Financial Studies* 26 (10): 2411–52.

List, John A. 2022. *The Voltage Effect: How to Make Good Ideas Great and Great Ideas Scale*. Currency.

Manso, Gustavo. 2016. “Experimentation and the Returns to Entrepreneurship.” *The Review of Financial Studies* 29 (9): 2319–40.

Powell, James L. 1986. “Censored Regression Quantiles.” *Journal of Econometrics* 32 (1): 143–55.

Ries, Eric. 2011. *The Lean Startup: How Today’s Entrepreneurs Use Continuous Innovation to Create Radically Successful Businesses*. Crown Business.

Scheier, Michael F., Charles S. Carver, and Michael W. Bridges. 1994. “Distinguishing Optimism from Neuroticism (and Trait Anxiety, Self-Mastery, and Self-Esteem): A Reevaluation of the Life Orientation Test.” *Journal of Personality and Social Psychology* 67 (6): 1063–78.

Scott, Erin L., Pian Shu, and Roman M. Lubynsky. 2019. “Entrepreneurial Uncertainty and Expert Evaluation: An Empirical Analysis.” *Management Science*, July, mnsc.2018.3244.

Stock, James H., and Motohiro Yogo. 2005. “Testing for Weak Instruments in Linear IV Regression.” In *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, 80–108. Cambridge University Press.

Tobin, James. 1958. “Estimation of Relationships for Limited Dependent Variables.” *Econometrica: Journal of the Econometric Society*, 24–36.

Vereshchagina, Galina, and Hugo A. Hopenhayn. 2009. “Risk Taking by Entrepreneurs.” *American Economic Review* 99 (5): 1808–30.

Wager, Stefan, and Susan Athey. 2018. “Estimation and Inference of Heterogeneous Treatment Effects Using Random Forests.” *Journal of the American Statistical Association* 113 (523): 1228–42.

Yadlowsky, Steve, Scott Fleming, Nigam Shah, Emma Brunskill, and Stefan Wager. 2021. “Evaluating Treatment Prioritization Rules via Rank-Weighted Average Treatment Effects.” *arXiv Preprint arXiv:2111.07966*.

FIGURES

Figure 1: Map of Locations of Incubator and Startup Locations

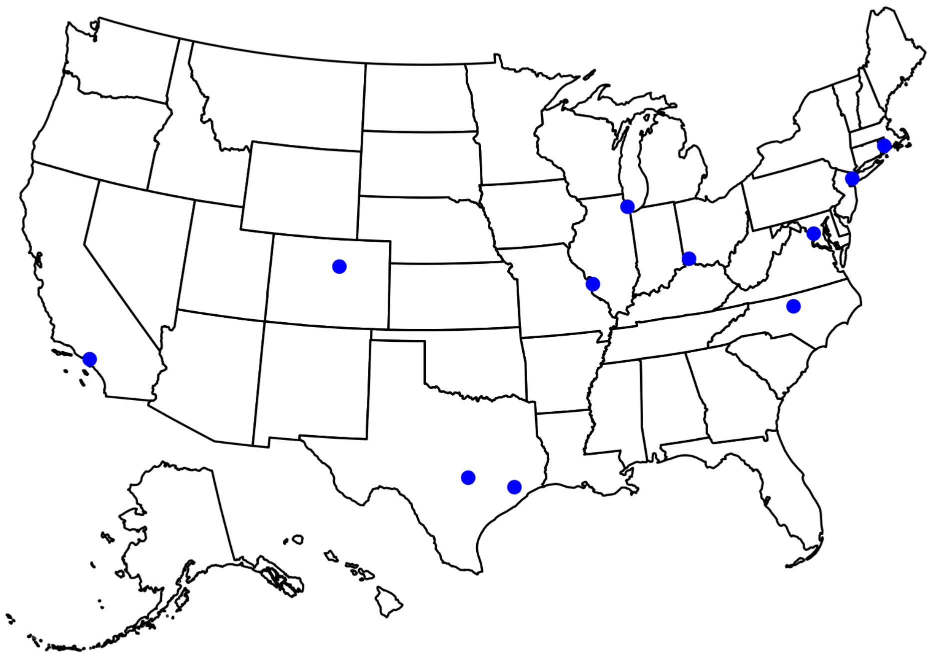

Note: This map shows the location of the incubators and startups enrolled in our field experiment. The twelve incubator locations inclusive of our pilot site (Matter in Chicago) are indicated by the blue dots.

Figure 2: Schematic of Experimental Design

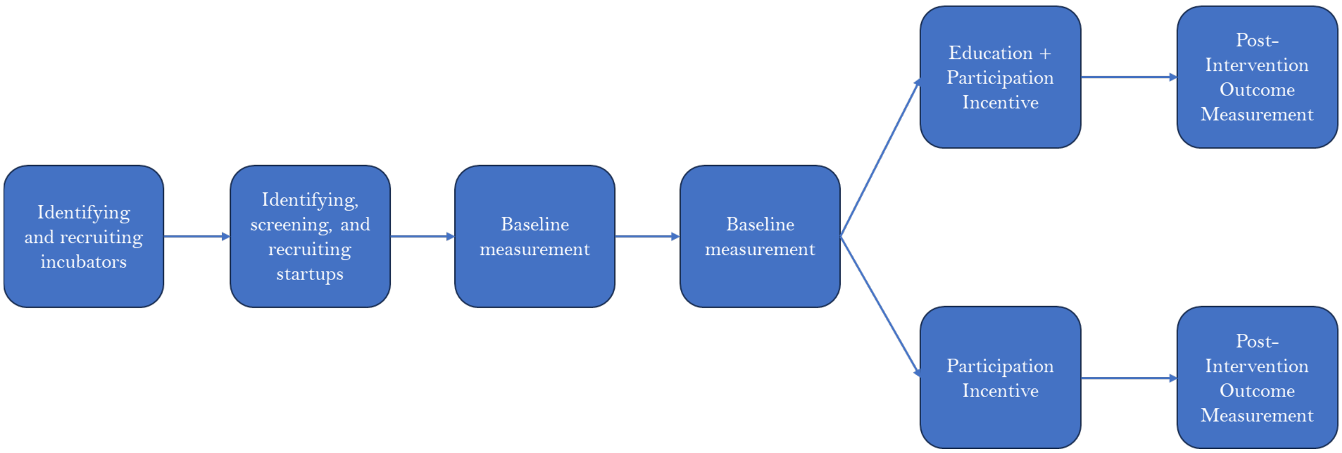

Figure 3: Kaplan-Meier Plots by Treatment Status

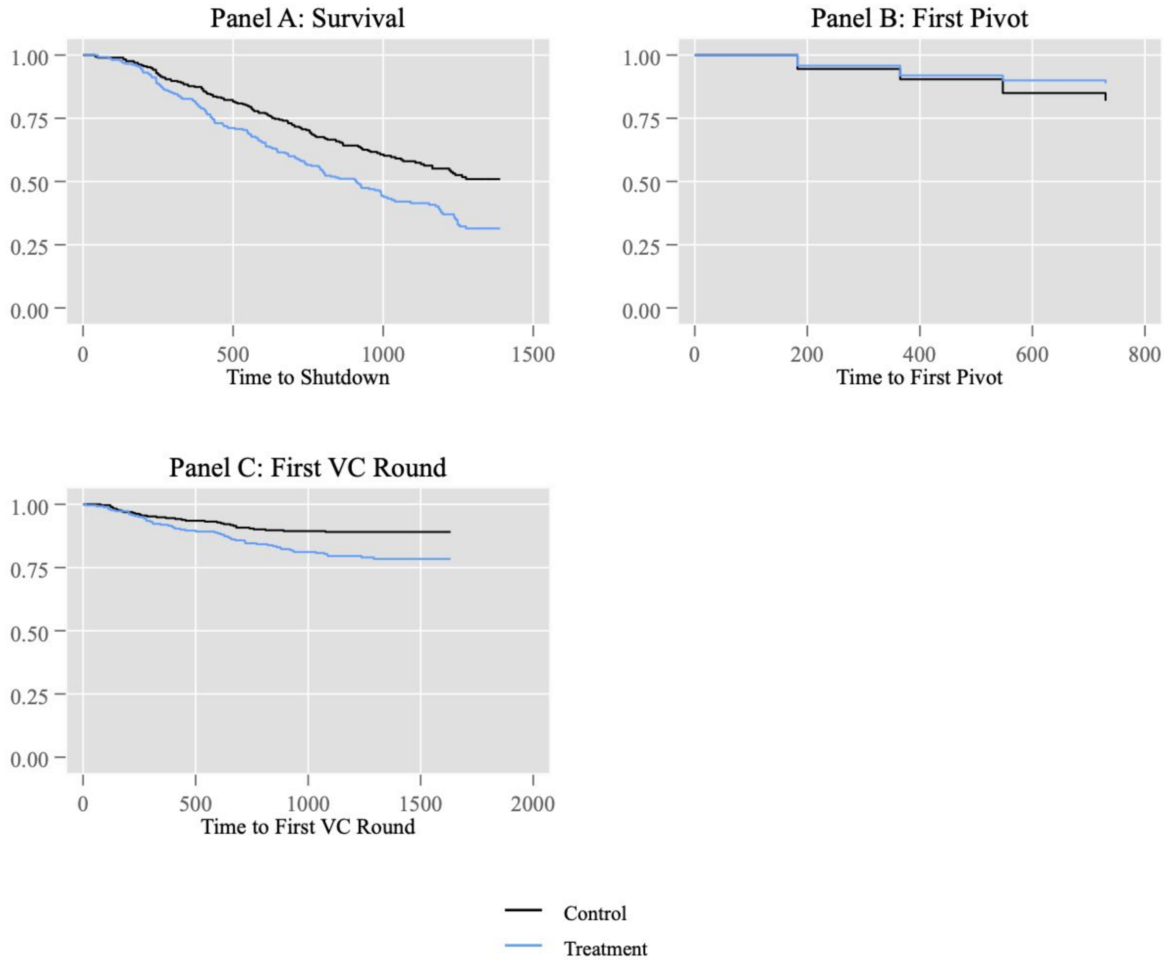

Notes. This figure contains a series of graphs that each contain plots of the Kaplan-Meier survivor function by treatment condition for different time-based dependent variables. In Panel A, we provide the plot of the Kaplan-Meier survivor function for time to shutdown for startups in the treatment and control conditions of our experiment. Log-rank test of the difference in time to shutdown between treated and control startups is statistically significant ($\chi^2 = 18.05$, $p < 0.000$). In Panel B, we show the plot of the Kaplan-Meier survivor function for time to first pivot for startups in the treatment and control conditions of our experiment. Log-rank test of the difference in time to first pivot between treated and control startups is statistically significant ($\chi^2 = 4.97$, $p < 0.025$). In Panel C, we show the plot of the Kaplan-Meier survivor function for time to first VC round for startups in the treatment and control conditions of our experiment. Log-rank test of the difference in time to first VC round between treated and control startups is statistically significant ($\chi^2 = 10.41$, $p < 0.001$). Log-rank test of the difference in time to first revenue between treated and control startups is not statistically significant ($\chi^2 = 0.22$, $p < 0.639$).

Figure 4: Kernel Density Plot of Funding by Treatment Status

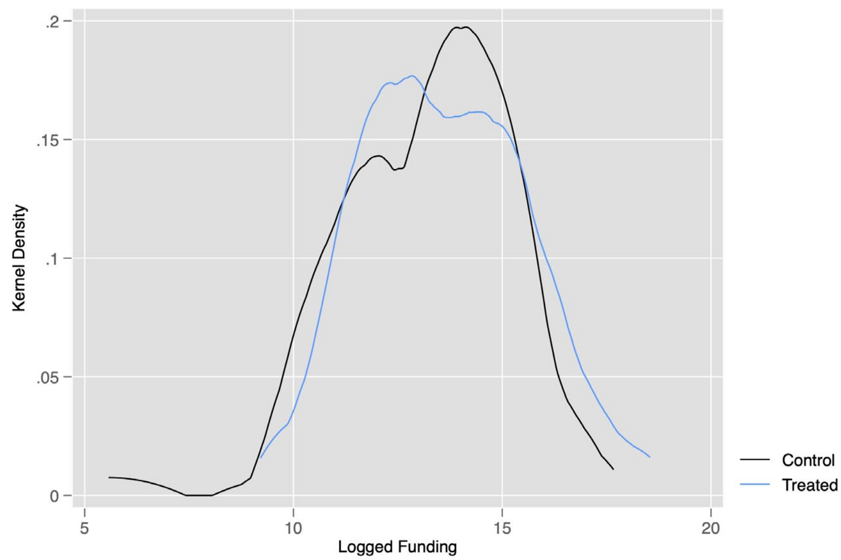

Notes. This figure provides kernel density plots of the distribution of Logged Funding for all non-zero observations of funding for treatment and control groups. Kolmogorov-Smirnov test of the equality of distributions provides evidence differences in the distribution of the treatment and control groups (p=0.061).

Figure 5: TOC Curves for Survival

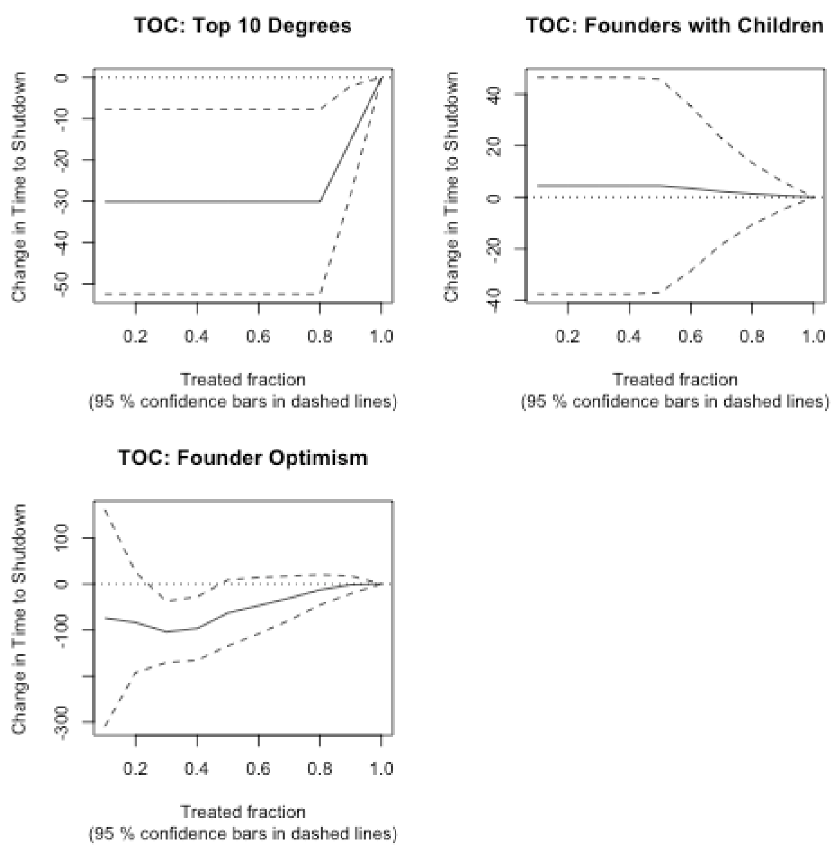

Figure 6: TOC Curves for Funding

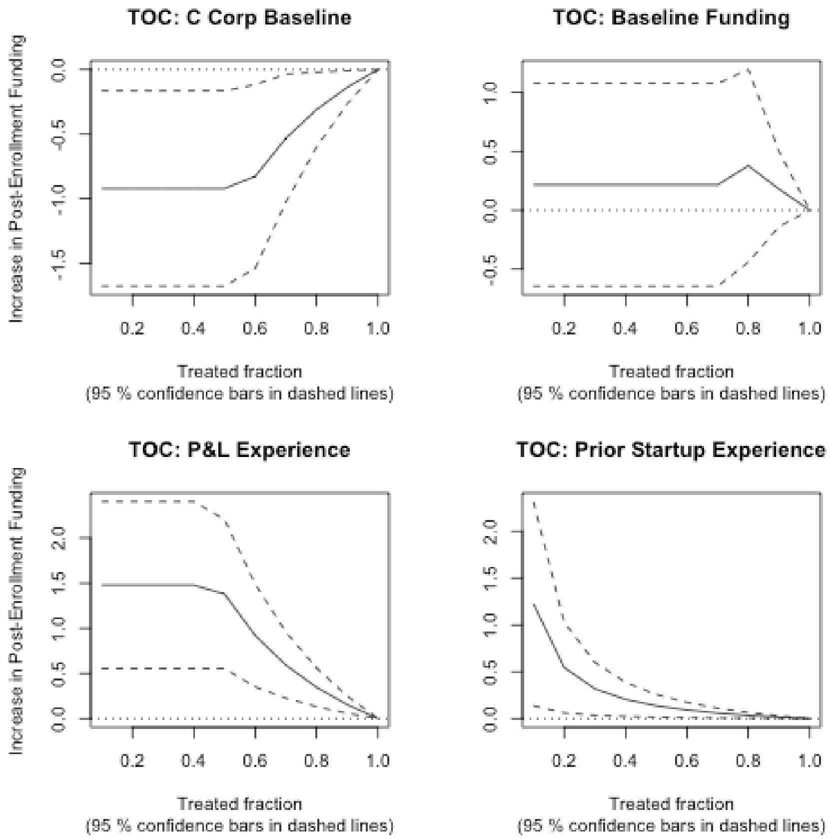

Figure 7: TOC Curve for Employment

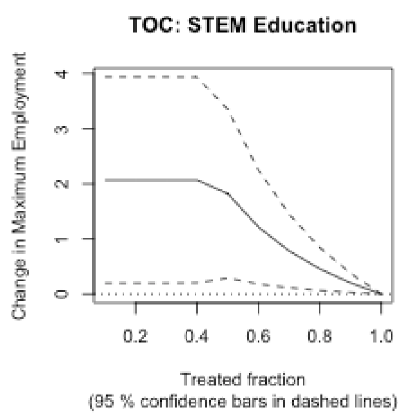

TABLES

Table 1: Sample Size by Location

Site	(1) Total	(2) Control	(3) Treatment
CIC - Boston	29	16	13
CIC - Cambridge	20	11	9
CIC - St. Louis	20	10	10
Capital Factory - Austin	60	31	29
1871 - Chicago	82	43	39
Union Hall - Cincinnati	44	25	19
Galvanize - Denver	34	18	16
Station - Houston	17	8	9
Cross Campus - Los Angeles	79	41	38
Matter - NYC	87	47	40
American Underground - Raleigh/Durham	48	26	22
1776 - Washington D.C.	33	17	16
Observations	553	293	260

Table 2: Survey Response Rates over Time

	Treatment Status		
	Control	Treated	Total
Survey			
Baseline	1.00	1.00	1.00
Follow-up 1 - 6 months	0.97	0.90	0.93
Follow-up 2 - 12 months	0.85	0.80	0.83
Follow-up 3 - 18 months	0.84	0.79	0.82
Follow-up 4 - 24 months	0.81	0.75	0.78

Table 3: Variable Names and Definitions

	Definitions
Outcome Variables	
Funding ($ 000,000)	Logged Total Funding Raised by Startup after enrollment
Time to Closure	Number of days from enrollment until the startup closes
Time to Pivot	Number of days from enrollment until the startup's first pivot
Time to Revenue	Number of days until the startup first generates revenue
Company Covariates	
Corporation	Was the startup a corporation (as opposed to LLC) at time of enrollment (1 if so, 0 otherwise).
Delaware Registration	Was the startup registered in Delaware (1 if so, 0 otherwise)
Baseline Funding	Total funding raised by startup before enrollment
Startup Founder Demographic Covariates	
Founder with Children	Does at least one founder have a child (1 if so, 0 otherwise)
BIPOC Founder	Based on survey response, does at least one founder come from an underrepresented group (1 if so, 0 otherwise)
Founder with P&L Experience	Based on survey response, does at least one founder have managerial experience with direct responsibility for profit and loss (P&L) of a department or division of a company
Married Founder	Is at least one founder married based on survey response (1 if so, 0 otherwise)
Female Founder	Based on survey response, is at least one founder female (1 if so, 0 othewise)
Prior Entre. Experience	Based on survey response, does at least one founder have experience working in a startup (1 if so, 0 otherwise)
STEM Education	Based on survey response, does at least one founder have a STEM graduate or undergraduate degree (1 if so, 0 otherwise)
Business Education	Based on survey response, does at least one founder have a graduate or undergraduate degree in business (1 if so, 0 otherwise)
Top 10 Degree	Based on survey response, does at least one founder hold a degree from a top ten university based on the US News and World Report Ranking in 2017
Top 11-20 Degree	Based on survey response, does at least one founder hold a degree from a university ranked 11-20 based on the US News and World Report Ranking in 2017
Average Age	Based on survey response, the average age of the founders of the startup at baseline.
Startup Founder Behavioral Covariates	
Max Founder Optimism	The highest level of founder-level optimism amongst the startup's founding team. Optimism is measured using the revised life orientation test (LOTR) described in Scheir et al. (1994)
Max Founder Time Preference	The highest level of founder-level time preference amongst the startup's founding team. Time preference is measured using the measure from Loewenstein, Read, and Baumeister (2003)
Max Founder Risk Preference	The highest level of founder-level risk preference amongst the startup's founding team. Risk preference is measured using the measure described in Barsky et al. (1997)

Notes. Additional information on variable construction is available in Online Appendix A

Table 4: Summary Statistics

	Observations	Mean	Median	Std. Dev.	Min	Max
Outcome Variables						
Logged Funding	553	3.59	0	5.85	0	17.67
Time to Closure	553	846.19	843	380.9	39	1390
Company Covariates						
Corporation	553	0.43	0	0.5	0	1
Delaware Registration	553	0.43	0	0.5	0	1
Baseline Funding	553	0.39	0	1.58	0	22
Startup Founder Demographic Covariates						
Founder with Children	553	0.46	0	0.5	0	1
BIPOC Founder	553	0.25	0	0.43	0	1
Founder with P&L Experience	553	0.53	1	0.5	0	1
Married Founder	553	0.56	1	0.5	0	1
Female Founder	553	0.37	0	0.48	0	1
Prior Entre. Experience	553	0.98	1	0.13	0	1
STEM Education	553	0.53	1	0.5	0	1
Business Education	553	0.46	0	0.5	0	1
Top 10 Degree	553	0.18	0	0.39	0	1
Top 11-20 Degree	553	0.11	0	0.31	0	1
Average Age	553	34.38	33.5	8.37	18	61
Startup Founder Behavioral Covariates						
Max Founder Optimism	553	19.95	21	3.48	6	24
Max Founder Time Preference	553	0.65	1	0.48	0	1
Max Founder Risk Preference	553	3.6	4	0.72	1	4

Table 5: Summary Statistics by Treatment Status and Test of Balance

	Treatment Status			
	Control	Treated	Total	Balance
Company Covariates				
Corporation	0.39	0.47	0.43	0.074*
	(0.49)	(0.5)	(0.5)	(0.043)
Delaware Registration	0.41	0.44	0.43	0.030
	(0.49)	(0.5)	(0.5)	(0.043)
Baseline Funding	0.29	0.5	0.39	0.022
	(1.05)	(2.01)	(1.58)	(0.013)
Startup Founder Demographic Covariates				
Female Founded	0.35	0.38	0.37	0.035
	(0.48)	(0.49)	(0.48)	(0.044)
Prior Entrepreneurship	0.98	0.98	0.98	-0.030
	(0.13)	(0.14)	(0.13)	(0.160)
STEM Education	0.56	0.5	0.53	-0.060
	(0.5)	(0.5)	(0.5)	(0.043)
Business Education	0.42	0.51	0.46	0.088**
	(0.49)	(0.5)	(0.5)	(0.042)
Top 10 Degree	0.17	0.19	0.18	0.036
	(0.38)	(0.39)	(0.39)	(0.055)
Top 11-20 Degree	0.11	0.11	0.11	-0.013
	(0.32)	(0.31)	(0.31)	(0.068)
Founder with P&L Experience	0.49	0.56	0.53	0.067
	(0.5)	(0.5)	(0.5)	(0.042)
Married Founder	0.55	0.58	0.56	0.032
	(0.5)	(0.49)	(0.5)	(0.043)
Founder with Children	0.45	0.47	0.46	0.023
	(0.5)	(0.5)	(0.5)	(0.043)
BIPOC Founder	0.23	0.26	0.25	0.040
	(0.42)	(0.44)	(0.43)	(0.049)
Average Age	34.03	34.77	34.38	0.003
	(8.19)	(8.56)	(8.37)	(0.003)
Startup Founder Behavioral Covariates				
Max Founder Optimism	20.04	19.85	19.95	-0.004
	(3.34)	(3.65)	(3.48)	(0.006)
Max Founder Time Preference	0.63	0.67	0.65	0.046
	(0.48)	(0.47)	(0.48)	(0.045)
Max Founder Risk Preference	3.58	3.63	3.6	0.024
	(0.76)	(0.67)	(0.72)	(0.030)

Notes. This table shows the distribution of the baseline covariates for the startups and their founders in our sample as well as the results of a balance test. The first column shows the mean and standard deviation (in parentheses) of the baseline covariates of the startups in the control group. The first column shows the mean and standard deviation (in parentheses) of the baseline covariates of the startups in the treatment group. The third column shows the mean and standard deviation (in parentheses) of the baseline covariates of all the startups in the sample. The fourth column shows the coefficient estimate of separate regressions for each covariate of the form $Treat = \beta_0 + \beta_1 \, Covar + \epsilon$. In the column, we show the estimates for each covariate of β_1 and its standard error (in parentheses). * indicates significance at the 10% level, ** at 5%, and ***at 1%.

Table 6: Raw Survival by Follow-up Period

Time Period	(1) Full Sample	(2) Control	(3) Treatment	(4) Diff. (3)-(2)	(5) P-Val
6 Months	96.38	96.93	95.77	-1.16	0.467
12 Months	85.35	87.71	82.69	-5.02	0.096
18 Months	75.77	80.55	70.38	-10.16	0.005
24 Months	65.64	71.67	58.85	-12.83	0.001
30 Months	48.28	52.90	43.08	-9.82	0.021
36 Months	30.20	34.13	25.77	-8.36	0.033
Observations	553	293	260	553	553

Notes. This table shows the raw differences in the percentage of the sample surviving during each of the 6-month periods in our sample as well as individual tests of the statistical significance of the difference at each 6-month interval. Column (1) shows the percentage of the total sample that survived in each 6-month follow-up period. Column (2) shows the survival percentage for the control group in each 6-month follow-up period. Column (3) shows the survival percentage for the treatment group in each 6-month follow-up period. Column (4) shows the difference in survival percentage between treatment and control groups. Column (5) shows the p-value associated with the t-statistic associated with a t-test of the difference in the distribution of surviving startups in the treatment and control groups in each 6-month follow-up period.

Table 7: Hazard Model of Startup Closure

	(1) Closure	(2) Closure	(3) Closure	(4) Closure
Treated	1.571***	1.585***	1.614***	1.594***
	(0.193)	(0.197)	(0.205)	(0.203)
Observations	553	553	553	553
Site Fixed Effects	X	X	X	X
Industry Fixed Effects	X	X	X	X
Company Covariates	-	X	X	X
Demographic Covariates	-	-	X	X
Behavioral Covariates	-	-	-	X

Notes. An observation in the sample represents the final status of the startup (closed or operating) and the survival time for each startup. Each column represents the output of a separate Cox proportional hazard regression predicting the hazard of closure and includes the key variable of interest, treatment, as well as site and industry fixed effects. Columns (2)-(4) include additional control covariates. Standard errors are clustered at the site level. * indicates significance at the 10% level, ** at 5%, and ***at 1%.

Table 8: Hazard Models of Time to First Pivot

Panel A: Cox Proportional Hazard

	(1) First Pivot	(2) First Pivot	(3) First Pivot	(4) First Pivot
Treated	0.670**	0.684*	0.698*	0.666*
	(0.125)	(0.137)	(0.150)	(0.152)
Observations	553	553	553	553
Site Fixed Effects	X	X	X	X
Industry Fixed Effects	X	X	X	X
Company Covariates	-	X	X	X
Demographic Covariates	-	-	X	X
Behavioral Covariates	-	-	-	X

Panel B: Competing Risk Model

	(1) First Pivot	(2) First Pivot	(3) First Pivot	(4) First Pivot
Treated	0.598*	0.605	0.604	0.545
	(0.184)	(0.190)	(0.205)	(0.205)
Observations	553	553	553	553
Site Fixed Effects	X	X	X	X
Industry Fixed Effects	X	X	X	X
Company Covariates	-	X	X	X
Demographic Covariates	-	-	X	X
Behavioral Covariates	-	-	-	X

Notes. In both panels of this table, an observation in the sample represents the pivot status of the startup (ever pivoted or not) and the time to first pivot for each startup. In Panel A, each column represents the output of a separate Cox proportional hazard regression predicting the hazard of pivoting and includes the key variable of interest, treatment, as well as site and industry fixed effects. Columns (2)-(4) include additional control covariates. Standard errors are clustered at the site level. In Panel B, each column represents the output of a separate Fine-Gray Competing Risk Models predicting the hazard of first pivot while controlling for the hazard of shutdown. * indicates significance at the 10% level, ** at 5%, and ***at 1%.

Table 9: Funding Outcomes by Treatment

Panel A: Models of Time to First VC Round

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Funding	Funding	Funding	Funding	Funding	Funding	Funding	Funding
Treated	2.134***	1.980***	2.091***	2.088***	1.778***	1.694**	1.813**	1.799**
	(0.382)	(0.410)	(0.579)	(0.566)	(0.377)	(0.395)	(0.549)	(0.527)
Observations	553	553	553	553	553	553	553	553
Site Fixed Effects	X	X	X	X	X	X	X	X
Industry Fixed Effects	X	X	X	X	X	X	X	X
Company Covariates	-	X	X	X	-	X	X	X
Demographic Covariates	-	-	X	X	-	-	X	X
Behavioral Covariates	-	-	-	X	-	-	-	X

Panel B: Funding Levels

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	I{Raise}	I{Raise}	I{Raise}	I{Raise}	Logged $	Logged $	Logged $	Logged $
Treated	0.108**	0.091**	0.085**	0.088**	1.412**	1.287**	1.219**	1.261**
	(0.035)	(0.031)	(0.033)	(0.033)	(0.515)	(0.419)	(0.457)	(0.448)
Observations	553	553	553	553	553	553	553	553
Site Fixed Effects	X	X	X	X	X	X	X	X
Industry Fixed Effects	X	X	X	X	X	X	X	X
Company Covariates	-	-	X	X	-	X	X	X
Demographic Covariates	-	-	X	X	-	-	X	X
Behavioral Covariates	-	-	-	X	-	-	-	X

Panel C: Cragg Hazard Models

	Logged Funding			
	(1)	(2)	(3)	(4)
Treated	0.555**	0.588***	0.697**	0.521
	(0.230)	(0.187)	(0.345)	(0.343)
Survival Time	0.001***	0.001*	0.001	0.000
	(0.000)	(0.000)	(0.000)	(0.000)
	I{Raise}			
	(1)	(2)	(3)	(4)
Treated	0.485***	0.455***	0.475***	0.474***
	(0.129)	(0.135)	(0.136)	(0.140)
Survival Time	0.001***	0.001***	0.001***	0.001***
	(0.000)	(0.000)	(0.000)	(0.000)
Observations	553	553	553	553
Site Fixed Effects	X	X	X	X
Industry Fixed Effects	X	X	X	X
Company Covariates	-	X	X	X
Demographic Covariates	-	-	X	X
Behavioral Covariates	-	-	-	X

Notes. This table presents multiple panels documenting the differences between treated and control groups in terms of time to funding and level of funding post-enrollment. In Panel A, an observation in the sample represents the VC funding status of the startup (ever funded or not) and the time to first VC funding round for each startup. Columns (1)-(4) show the output of separate Cox proportional hazard regressions predicting the hazard of receiving a first VC funding round and includes the treatment variable, site and industry fixed effects, and covariates depending on the regression. Columns (5)-(8) show the output of Fine-Gray competing risk models predicting the hazard of a first VC funding round while controlling for the hazard of shutdown. In Panel B, an observation is level of funding at endline. In columns (1)-(4) we show the results of linear probability models predicting if the startup ever received external funding. In columns (5)-(8) we show the results of OLS regressions of treatment on logged funding. All models include site and industry fixed effects as well as covariates described in the panel. In Panel C, an observation is level of funding at endline. Each column shows the output of Cragg hurdle regression models that first predict whether a startup received funding and then the level of logged funding conditional on raising. All models include site and industry fixed effects as well as covariates described in the panel. * indicates significance at the 10% level, ** at 5%, and ***at 1%.

Table 10: Maximum Employment in Startup

	Maximum Employment			
	(1)	(2)	(3)	(4)
Treated	1.286*	1.275*	1.218**	1.218**
	(0.169)	(0.186)	(0.110)	(0.118)
Observations	553	553	553	553
Site Fixed Effects	X	X	X	X
Industry Fixed Effects	X	X	X	X
Company Covariates	-	X	X	X
Demographic Covariates	-	-	X	X
Behavioral Covariates	-	-	-	X

Notes. This table examines the impact of treatment on the maximum level of employment that the startup reached during the follow-up period. Each column represents the output of a separate Poisson regression with standard errors clustered at the site level. The coefficients are exponentiated so that they represent the percentage difference of treatment relative to control where a coefficient of 1.2 represents a 20% increase in the level of employment for the treated group relative to the control group. * indicates significance at the 10% level, ** at 5%, and ***at 1%.

Table 11: Entrepreneurship Choice in Next Job Role

	Founded Startup Next			
	(1)	(2)	(3)	(4)
Treated	-0.092***	-0.090**	-0.110***	-0.110***
	(0.029)	(0.030)	(0.032)	(0.035)
Observations	434	434	434	434
Site Fixed Effects	X	X	X	X
Industry Fixed Effects	X	X	X	X
Company Covariates	-	-	X	X
Demographic Covariates	-	-	X	X
Behavioral Covariates	-	-	-	X

Notes. The sample for this table is all founders of the startup in our sample that had a LinkedIn profile we could use to observe their choice of entrepreneurship after shutdown of the startup enrolled in our study. An observation in this sample represents whether or not a founder for one of the startups in our sample founded another startup directly after shutting down the enrolled startup. All regressions are limited probability models predicting whether a startup founder chooses to create another startup immediately after shutting down their startup. Standard errors clustered at the site level. * indicates significance at the 10% level, ** at 5%, and ***at 1%.

Table 12: Treatment on Treated

Panel A: First Stage

	Sessions Attended		Attended Any Sessions	
	(1)	(2)	(3)	(4)
Treated	2.163***	2.201***	0.653***	0.651***
	(0.155)	(0.154)	(0.038)	(0.037)
Average First-Stage F-Statistic	216.4	231.23	273.78	298.22
Observations	553	553	553	553
Site Fixed Effects	X	X	X	X
Industry Fixed Effects	X	X	X	X
Company Covariates	-	X	-	X
Demographic Covariates	-	X	-	X
Behavioral Covariates	-	X	-	X

Panel B: IV Estimates

	Time to Shutdown			
	(1)	(2)	(3)	(4)
Sessions Attended	-47.585**	-48.350***		
	(19.060)	(17.761)		
Attended Any Sessions			-157.716**	-162.796***
			(64.833)	(59.418)
	I{Raise}			
	(1)	(2)	(3)	(4)
Sessions Attended	0.050***	0.040***		
	(0.017)	(0.015)		
Attended Any Sessions			0.166***	0.135***
			(0.055)	(0.050)
	Logged Funding			
	(1)	(2)	(3)	(4)
Sessions Attended	0.709***	0.573***		
	(0.227)	(0.202)		
Attended Any Sessions			2.351***	1.936***
			(0.748)	(0.675)
	Maximum Employment			
	(1)	(2)	(3)	(4)
Sessions Attended	1.135**	1.075*		
	(0.066)	(0.047)		
Attended Any Sessions			1.458**	1.263*
			(0.235)	(0.173)
	Founded Startup Next			
	(1)	(2)	(3)	(4)
Sessions Attended	-0.046***	-0.056***		
	(0.016)	(0.017)		
Attended Any Sessions			-0.143***	-0.176***
			(0.047)	(0.052)
Site Fixed Effects	X	X	X	X
Industry Fixed Effects	X	X	X	X
Company Covariates	-	X	-	X
Demographic Covariates	-	X	-	X
Behavioral Covariates	-	X	-	X

Notes. This table seeks to examine the impact of session attendance on key dependent variables by using assignment to treatment as an instrument for attendance of sessions. Panel A shows the first stage of the IV regressions. Panel B shows the results of a series of IV regressions showing the impact of attendance measured multiple ways on performance. In Table B3 of Online Appendix B, we also show the robustness of our Logged Funding results to an alternative transformation using the inverse hyper sine function. * indicates significance at the 10% level, ** at 5%, and ***at 1%.

Table 13: Best Linear Predictor of Covariates on Estimated CATE

Variable	(1) Survival	(2) Funding	(3) Employment
Corporation	110.713	3.423***	-0.588
	(82.727)	(0.827)	(3.233)
Delaware Registration	-108.109	-1.359	2.763
	(113.466)	(1.691)	(4.315)
Baseline Funding	-18.403	-0.911***	0.810
	(32.795)	(0.281)	(1.800)
Age	4.111	-0.045	0.223
	(4.767)	(0.070)	(0.137)
Female Founded	-51.591	1.693	0.648
	(91.607)	(1.324)	(2.493)
Prior Entrepreneurship	-1.143	-10.109*	-32.447
	(478.371)	(5.792)	(42.411)
STEM Education	27.268	2.018	-3.917**
	(86.427)	(1.274)	(1.895)
Business Education	108.191	0.319	0.348
	(68.388)	(1.425)	(3.549)
Top 10 Degree	162.021*	2.131	1.094
	(86.009)	(1.632)	(3.163)
Top 11-20 Degree	33.174	0.859	-1.515
	(99.975)	(2.336)	(2.916)
Founder with P&L Experience	-93.057	-2.694***	0.533
	(77.554)	(1.021)	(1.484)
Married Founder	156.606	0.233	0.785
	(106.553)	(1.135)	(4.068)
Founder with Children	-146.947**	-0.019	0.440
	(59.267)	(1.423)	(2.866)
BIPOC Founder	18.547	0.490	-1.867
	(110.650)	(1.548)	(2.511)
Max Founder Optimism	19.696*	0.159	-0.396
	(11.198)	(0.154)	(0.341)
Max Founder Time Preference	-69.711	-0.810	-1.703
	(87.322)	(0.860)	(2.746)
Max Founder Risk Preference	51.465	0.878	0.441
	(61.031)	(0.553)	(1.880)
Observations	553	553	553

Notes. This table explores potential heterogeneity by measuring the relationship between the predicted conditional average treatment effect (CATE) from our calibrated generalized random forest models (GRF) and the baseline covariates of the startup. To do so, we use a best linear prediction of the covariates onto the distribution of our predicted CATE. The results in the two columns correspond to the GRF models for the two main outcome variables of our study, funding and survival and their standard errors. * indicates significance at the 10% level, ** at 5%, and ***at 1%.