



Predicting entrepreneurial success is hard: Evidence from a business plan competition in Nigeria[☆]

David McKenzie^{a,*}, Dario Sansone^{b,c}

^a Development Research Group, The World Bank Group, USA

^b Georgetown University, USA

^c University of Exeter, UK



ARTICLE INFO

JEL classification:

C53

L26

M13

Keywords:

Entrepreneurship

Machine learning

Business plans

Nigeria

ABSTRACT

We compare the absolute and relative performance of three approaches to predicting outcomes for entrants in a business plan competition in Nigeria: Business plan scores from judges, simple ad-hoc prediction models used by researchers, and machine learning approaches. We find that i) business plan scores from judges are uncorrelated with business survival, employment, sales, or profits three years later; ii) a few key characteristics of entrepreneurs such as gender, age, ability, and business sector do have some predictive power for future outcomes; iii) modern machine learning methods do not offer noticeable improvements; iv) the overall predictive power of all approaches is very low, highlighting the fundamental difficulty of picking competition winners.

1. Introduction

Millions of small businesses are started every year in developing countries. However, more than a quarter of these die within their first year (McKenzie and Paffhausen, 2018), while only a small subset of firms grows rapidly, creating disproportionate value in terms of employment and incomes (Olafsen and Cook, 2016). The ability to identify ex ante which firms will succeed is of key interest to investors seeking to maximize returns, determining the extent to which capital can be allocated to the highest return projects. Being able to identify these high growth potential firms is also important for governments seeking to target programs to these firms (OECD, 2010), and is of interest to researchers seeking to characterize what makes a successful entrepreneur. Moreover, if the characteristics that are predictive of high growth are malleable (and additional studies show a causal link between such characteristics and firm outcomes), this research can also spur policy efforts to attempt to change these attributes in individuals lacking them.

Business plan competitions are increasingly used in both developed

and developing countries to attempt to spur high growth entrepreneurship. These competitions typically attract contestants with growth prospects that are much higher than the average firm in the economy. But then the key question is whether one can predict which firms from among these applicants have better growth prospects. We use data from applicants to Nigeria's YouWiN! program, the world's largest business plan competition (Kopf, 2015), to answer this question. We compare the absolute and relative performance of three different methods for identifying which firms will be most successful: the standard approach of relying on expert judges, simple ad hoc prediction models used by researchers that employ a small number of variables collected in a baseline survey, and modern machine learning (ML) methods that consider 566 possible predictors and non-linear ways of combining them.

Business plan competitions typically rely on expert judges to score proposals, with higher scores given to those businesses judges view as having higher likelihoods of success. This approach has the advantage of using context-specific knowledge and enabling holistic judgement of business plans. However, using judges to score proposals can be costly

[☆] We are grateful to the editor Jeremy Magruder, two anonymous referees, Eric Auerbach, Keisuke Hirano, William Maloney, Givi Melkadze, and Allison Stashko for their helpful comments. Participants at the CSAE, AI and Development, and WADES conferences and seminar participants at Georgetown University, The World Bank and UIUC provided useful feedback. Funding from the Strategic Research Program (SRP) of the World Bank is gratefully acknowledged. An earlier version of this paper has circulated under the title "Man vs. Machine in Predicting Successful Entrepreneurs: Evidence from a Business Plan Competition in Nigeria".

* Corresponding author.

E-mail addresses: dmckenzie@worldbank.org (D. McKenzie), ds1289@georgetown.edu (D. Sansone).

<https://doi.org/10.1016/j.jdevec.2019.07.002>

Received 25 April 2018; Received in revised form 3 July 2019; Accepted 10 July 2019

Available online 24 July 2019

0304-3878/© 2019 Published by Elsevier B.V.

and time-consuming, and there are also concerns that human evaluations can sometimes lead to discriminatory decisions (Blanchflower et al., 2003), or be driven by overconfidence (Zacharakis and Shepherd, 2001). One alternative may then be to attempt to predict which businesses will succeed based on simple regression models that include survey variables like gender, age, and ability suggested by theory and researcher judgement (Fafchamps and Woodruff, 2017; McKenzie, 2015). However, the selection of which variables to include in these models is ad hoc and, since it reflects a researcher's opinion, may omit important determinants of future success. Machine learning techniques offer an alternative approach that is less reliant on human judgement, and are designed specifically for out-of-sample prediction (Kleinberg et al., 2015). They offer the potential for better prediction performance by using high-dimensional data reduction methods to identify patterns that might be too subtle to be detected by human observations (Luca et al., 2016). This could then offer a cheaper and more effective alternative to identifying which businesses will succeed.

Using data on more than 1050 winners and more than 1050 non-winners, we find that the business plan scores from judges do not significantly predict survival, or employment, sales and profits outcomes three years after entry. This is not due to the scores all being similar to one another, nor to the outcomes being similar for all firms. For example, among the non-winners, a firm at the 90th percentile has 13 times the employment and 2.4 times the profits of a firm at the 10th percentile. Rather, conditional on getting through the first phase of the competition, the scores of judges do not predict which firms will be more successful.

We then use several ad hoc prediction models. One set of these models are simple univariate prediction models using variables like gender, age, education, and ability that the literature has suggested might be related to firm performance. A second set of models comes from a simple logit model based on this literature, used by McKenzie (2015), and from a similar model used by Fafchamps and Woodruff (2017). We do find that some observable characteristics of entrepreneurs help predict future success, with high-ability males in their thirties doing better. Nevertheless, the overall accuracy of these models is still low.

These results are then contrasted with out-of-sample predictions from three modern ML approaches - LASSO, Support Vector Machines, and Boosted Regression - that use data reduction techniques and highly nonlinear functions to make predictions starting with 566 possible predictors. These methods may end up choosing interactions of variables and survey responses that would be unlikely to be used by human experts. However, we find that the overall accuracy of these models is similar to that of the simpler ad hoc models, and thus also low. Although we find that machine learning is often unable to beat simple predictors or simple models from economists when it comes to predicting average performance, we do find some role for machine learning and simple models in identifying the top tail of high-growth firms. Nevertheless, even our best models are typically only able to ex ante identify only 20 percent of the firms that will end up in the top decile of the outcome distribution.

We then examine different potential explanations for why the judges and machine-learning algorithms do not perform better. We discuss whether judges were trying to predict a different outcome, whether the use of anonymized plans restricted their ability to statistically discriminate, whether they lacked population-specific knowledge, and whether they are better at telling apart the bottom tail. None of these appear to be the main reasons for their poor performance. We then examine whether the machine-learning performance is hampered by insufficient sample, by a lack of independent variation in the inputs, or by poor model choice. Again, these do not seem to be the main factors behind the poor performance.

Taken together, these results point to the fundamental difficulty of identifying in advance which entrepreneurs will be more successful from

among a sample that has already made it through a first phase of a business plan competition. This is a feature of the fundamental riskiness inherent in entrepreneurship (Hall and Woodward, 2010), and it is consistent with the inability of professional investors to know in advance whether a technology or venture will succeed, conditional on it passing some initial screening (Kerr et al., 2014; Hall and Woodward, 2010).

This paper builds on two existing literatures. The first is a literature on predicting which firms will succeed. Most of the existing literature looks at start-ups and venture-backed firms in developed countries, and while some of these studies find that judges' scores have some predictive power (e.g. Astebro and Elhedhli, 2006; Scott et al., 2015), they also point to the immense difficulty of identifying who will be more successful (Kerr et al., 2014; Nanda, 2016). Data analysts have also started to predict which new tech firms will be successful by looking at which sub-sectors have attracted the highest investments by venture capitalists and which companies have raised funds from the top-performing investment funds in the world (Ricadela, 2009; Reddy, 2017). There has been much less study of this issue in developing countries (Nikolova et al., 2011). An important exception is Fafchamps and Woodruff (2017), who compare the performance of judges' evaluations to survey-based measures, and find that both have some (independent) predictive power among business plan contestants in Ghana. We build on their work with a much larger sample, and through the incorporation of machine learning.

Secondly, we contribute to a growing literature that uses machine learning in economics (Mullainathan and Spiess, 2017). Kleinberg et al. (2015) note that this method may be particularly useful for a number of policy-relevant issues that are prediction problems, rather than causal inference questions. Economists have started applying these methods to a number of different settings such as predicting denial in home mortgage applications (Varian, 2014), quality of police and teachers (Chalfin et al., 2016), poverty from satellite data (Jean et al., 2016), conflicts in developing countries (Celiku and Kraay, 2017), if defendants will commit crimes when released on bail (Kleinberg et al., 2017), and high school and college dropout (Aulck et al., 2016; Sansone, 2018). Our work shows the limits of this approach when it comes to predicting entrepreneurial success, and that machine learning need not necessarily improve on the performance of simpler models that have been more commonly used by economists.

2. Context and data

Business plan competitions are a common approach for attempting to attract and select entrepreneurs with promising business ideas, and have been increasingly used in developing countries as part of government efforts to support entrepreneurship. The typical competition attracts applicants with the promise of funding, training or mentoring, and/or publicity for the winners. Applications are then usually subjected to an initial screening, with the most promising plans then being scored and prizes awarded to the top-scoring entries.

2.1. The YouWiN! competition

Our context is the largest of these competitions, the *Youth Enterprise with Innovation in Nigeria* (YouWiN!) competition. This competition was open to all Nigerians aged 40 or younger, who could apply to create a new firm or expand an existing one. We work with the first year of the program, which had a closing date for initial applications of November 25, 2011. The competition attracted 23,844 applications. A first stage basic scoring selected the top 6000 applicants, who were offered a 4-day business plan training course. Those who attended the course were then asked to submit a detailed business plan, with 4517 proposals received. These proposals were then scored by judges - in a process described below - with the highest scores used to select 2400 semi-finalists. 1200

among those were then chosen as winners and received grants averaging almost US\$50,000 each. 480 of these winners were chosen on the basis of having the top scores nationwide or within their region, while 720 winners were randomly chosen from the remaining semi-finalists.¹

McKenzie (2017) uses this random allocation to measure the impact of winning this competition, and provides more details on this competition and selection process. His analysis found that being randomly selected as a competition winner resulted in greater firm entry from those with new ideas, higher survival of those with existing businesses, and higher sales, profits, and employment in their firms. It shows that the competition was able to attract firms that, on average, had high growth prospects. Our goal in this paper is to use this setting to see whether the judge scores, simple heuristic models, or machine learning approaches can identify *ex ante* which firms from among these applicants will be most likely to survive and grow.

2.2. Data

Our starting sample consists of 2506 firms that are comprised of 475 non-experimental winners (the national and regional winners), 729 experimental winners, 1103 semi-finalists that are in the control group, and 199 firms that submitted business plans but did not score high enough for the semi-finals.² By considering non-winners in addition to winners, we are able to test how well human experts and machine learning approaches do in forecasting firm growth for both a sample of firms that receive financial support, and for a sample of firms that are not receiving grant support. The majority of this sample consists of individuals seeking to create new firms (63%), with 37% coming from existing firms looking to expand. For each of these firms we have data from their initial online application, from the submission of their business plans, the score provided by judges on their business plans, and short-term follow-up survey data on some ability and personality measures that are assumed to be stable over time. We discuss each in turn, and provide a full description of the different variables in [Appendix A.1](#).

The application form required individuals to provide personal background such as gender, age, education, and location, along with written answers to questions like “why do you want to be an entrepreneur?” and “how did you get your business idea?”. We construct several simple measures of the quality of these written responses, such as the length of answer required, whether the application is written in all capital letters, and whether they claim to face no competition. We also use administrative data on the timing of when applications were submitted relative to the deadline.

The business plan submission contained more detailed information, including a short baseline survey instrument. This collected additional information about the background of the entrepreneur, including marital status, family composition, languages spoken, whether they had received specific types of business training, employment history, time spent abroad, and ownership of different household assets. It elicited risk preferences, asked about motivations for wanting to operate a business, and measured self-confidence in the ability to carry out different

¹ Selection of winners was stratified by whether the proposal was for an existing or new firm, and by region. 300 National winners (225 existing, 75 new) were selected as the top overall, then 30 regional winners were chosen from each of the six regions from a shortlist of 45 existing firms and 15 new firms per region. The 720 experimental winners were chosen as 120 per region, again oversampling existing firm applicants. See McKenzie (2017) for additional details.

² McKenzie (2017) dropped 5 non-experimental winners who were disqualified by the competition, and includes 9 experimental winners who were disqualified. We drop these 14 firms, along with 9 control group firms that were non-randomly chosen as winners. In addition, McKenzie (2017) attempted to survey 624 firms that were not selected for the business plan training. Since these firms did not submit business plans, they do not have business plan scores and are thus excluded from our analysis.

entrepreneurial tasks. More detailed information about the business they wanted to operate or their existing business they intended to expand included the name of the business, business sector, formal status, access to finance, taxes paid, challenges facing the business, and projected outcomes for turnover and employment. Unfortunately, we do not have access to the written text of the business plans, and so cannot employ textual analysis of the plans themselves.

The initial application and business plan submission were all done online. While they contain a rich set of data on these applicants, there were some measures which needed to be collected face-to-face with explanations, and so were left for follow-up surveys. Two ability measures were collected in the follow-up surveys: digit-span recall, and a Raven progressive matrices test of abstract reasoning. In addition, the grit measure of Duckworth et al. (2007) was also collected during these short-term follow-ups. We assume that these measures are stable over one to two years, and therefore include them as potential characteristics that could be used to predict future outcomes had they been collected at the time of application. To be consistent with the set of predictors included in Fafchamps and Woodruff (2017), we also include two measures of current and expected life satisfaction, as well as additional measures of self-confidence.

2.3. Measuring success

McKenzie (2017) tracked these firms in three annual follow-up surveys, and then in a five-year follow-up during a period of recession in Nigeria. We use data from the three-year follow-up survey, fielded between September 2014 and February 2015, to measure business success in this paper. This survey had a high response rate, with data on operating status and employment available for 2128 of the 2492 firms in our sample of interest (85.4%).³ We believe this is a more useful test than the five-year follow-up, which, in addition to a lower response rate, occurs in an atypical period in which Nigeria suffered its worst economic performance in thirty years, resulting in many firms struggling. Moreover, three years is the time horizon over which business plans were prepared (see next section).

We focus on four metrics of success: business operation and survival, total employment, sales, and profits. [Appendix A.1](#) describes how these are measured. 93.6% of the competition winners were operating businesses three years after applying, compared to 58.4% of the control group and other non-winners. [Fig. 1](#) then plots the distributions of employment, sales, and profits outcomes for the winners and non-winners. We see substantial dispersion in the outcomes of these firms three years after they applied for the program. For example, among winning firms in operation, a firm at the 90th percentile has 20 employees, five times the level of a firm at the 10th percentile of employment; and a firm at the 90th percentile of profits earns 500,000 naira per month (USD \$2730), around four times the level of a firm at the median, and 100 times that of a firm at the 10th percentile. Likewise, while average levels of these outcomes are lower for non-winners, there is also a long right tail for this group, with some firms doing substantially better than others. Given how much variability there is in these outcomes, it is therefore of interest to examine whether the more successful of these firms could have been identified in advance.

³ [Appendix 8](#) of McKenzie (2017) examines the characteristics of this attrition. It shows that attrition rates were 6% higher for non-winners than winners, but that this did not change the balance on baseline observables (including business scores). Our analysis is done separately for winners and non-winners, and we cannot reject that round 3 attrition is orthogonal to the business plan score for any of our sample groups. Coupled with the relatively low absolute level of attrition, this suggests that attrition is unlikely to drive our results, and we thus focus on the sample of firms interviewed in this three-year follow-up.

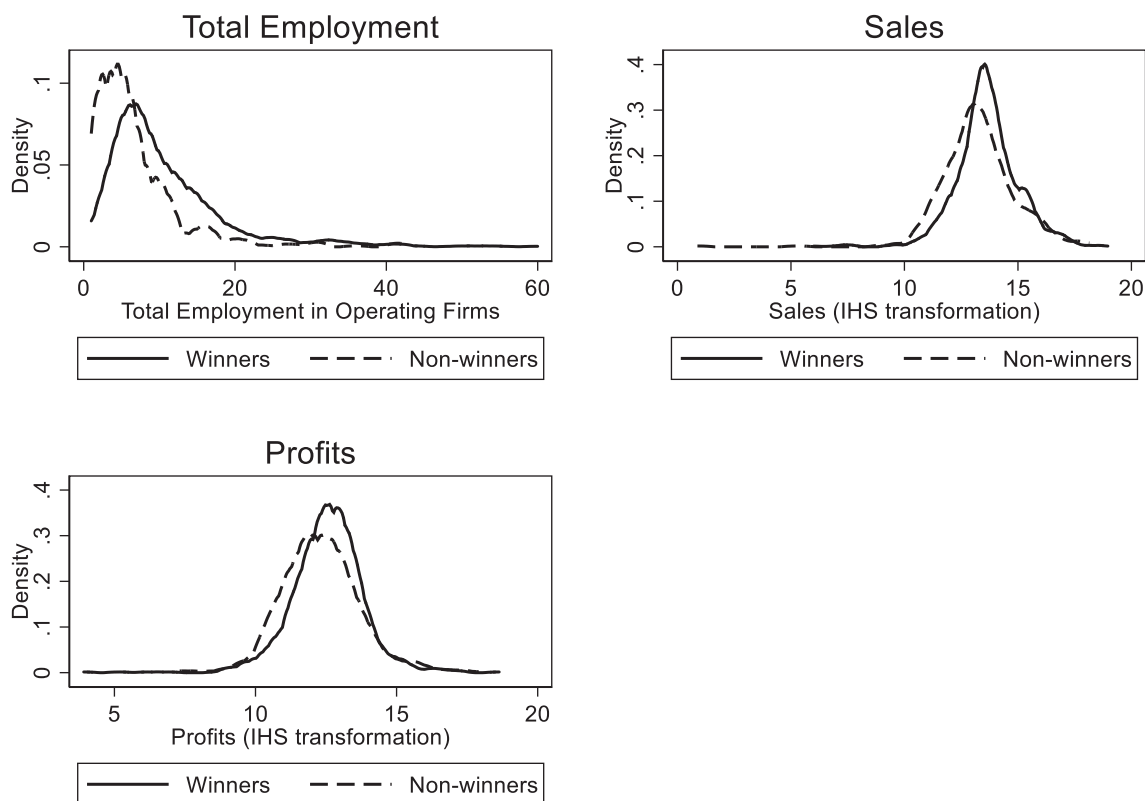


Fig. 1. Substantial Dispersion in Firm Outcomes Three Years After Applying. Notes: Figure shows kernel densities of firm outcomes for firms surveyed in third follow-up, and are conditional on the firm operating at that time. Winners include both non-experimental and experimental winners, and non-winners include control group semi-finalists as well as a sample of firms submitting business plans which had first round application scores just above the cut-off for getting selected for business plan training. Total employment truncated at 60 (9 outliers) for readability. Profits and Sales are measured in monthly Naira, and have been transformed using the inverse hyperbolic sine transformation.

3. Human approaches to predicting business success

The standard approach to choosing businesses to support is to rely on the human judgement of experts. We consider prediction by two types of experts. The first group of experts is composed by the judges in the business plan competition, while the second group is composed by economists who generate their prediction models based on their expert judgements of what variables should predict business success.

3.1. Using judges to evaluate business plans

The business plans prepared by applicants were lengthy documents, averaging 13,000 words of narrative coupled with spreadsheets of financial projections. They contained a detailed description of the proposed business, market and industry analysis, production plans, operations and logistic details, customer analysis, financial analysis, competitive environment, and other details, along with financing plans, discussion of what the grant would be used for, discussion of the risks, threats, and opportunities, and sales and growth projections for the next three years.

Scoring for the competition was overseen by an independent project unit supported by DFID and the World Bank, with significant effort taken to ensure that the scoring was impartial and rigorous. Scoring was done by the Enterprise Development Center (EDC), a sister unit of the Lagos Business School within the Pan-Atlantic University in Lagos; and by the Nigerian office of PriceWaterhouseCoopers (PwC). Each selected 10 judges, who were experts in business and entrepreneurship and knowledgeable about local conditions. These judges included business owners and alumni of Lagos Business School, as well as experts involved in working directly with successful businesses in Nigeria. These judges were

given training, which included marking sample plans and receiving feedback. Markers had also access to a forum where they could seek advice from their peers and resolve issues.

In order to ensure impartiality, and to be able to score such a large number of applications, applicants did not make in-person pitches to judges, but were scored only on the basis of their business plan. These plans were anonymized. Judges were given a set template for scoring, designed by the two leads from EDC and PwC, which assigned marks on 10 different criteria covering the management skills and background of the owner, the business idea and market, financial sustainability and viability, job creation potential, and a capstone score to allow for their overall assessment of the chance of business success (Appendix A.1). Importantly, these judges were asked to assess which firms were most likely to succeed, not which firms would have the largest treatment effects from receiving a grant from the competition. In particular, judges were told that they were looking for submissions that “stand out in their innovativeness, feasibility and job creation potential”. A typical plan took 30–45 min to mark. A random sample of 204 of the business plans were then re-scored by an independent team from Plymouth Business School and any significant disparities reviewed. The mean score in this sample was within 0.4 points out of 100 of the mean from the original judges, with this difference not statistically significant, and their quality assurance concluded that “the marking criteria deployed by LBS/EDC in this case did offer a broad mechanism to differentiate leading business plans from less convincing business plans”.

Fig. 2 plots the distribution of business plan scores according to their competition outcome. Note that the distribution for non-experimental winners overlaps with that of the experimental winners and losers because of the preference given towards existing firm applicants, and because of differences across regions in the scoring threshold needed to

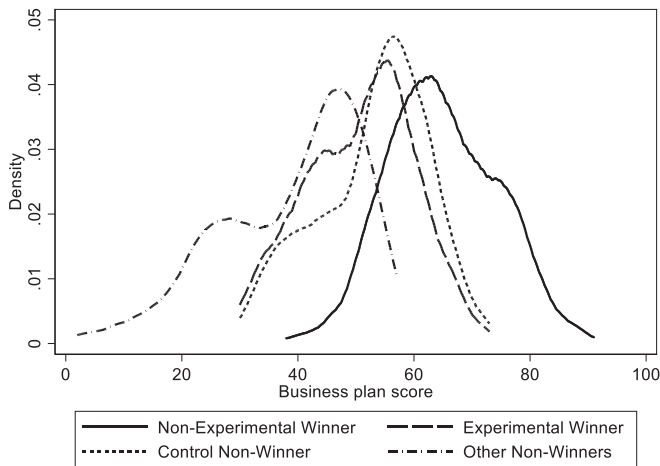


Fig. 2. Distribution of Business Plan Scores by Competition Status. Notes: kernel densities of business plan scores shown. Non-experimental winners were chosen on the basis of having the top scores nationwide, or as being the most promising within their geographic region, with fixed quotas by whether the application was for a new or existing business. Experimental winners and control non-winners were chosen randomly from within the remaining semi-finalists, with this randomization done within strata defined by region and existing or new business status. Other non-winners are the scores for our remaining sample of non-winners, selected on the basis of only just making the first round cut to receive business plan training.

be a regional winner. We see that distributions of scores are wide for both the winners and the non-winners in our sample. The winning scores range from 30 to 91 out of 100, with a mean of 56 and standard deviation of 12. The scores of the non-winners in our sample range from 2 to 73 out

of 100, with a mean of 51 and a standard deviation of 11. This shows that the judges did view firms as substantially different from one another in terms of their potential for business success. Moreover, we have seen that the firms also had wide dispersion in subsequent outcomes. This then raises the question of whether the judges were able to identify in advance which firms were more likely to succeed.

3.2. Do judges' scores predict business success?

Since the winning firms received substantial business grants, which were shown in McKenzie (2017) to have a causal impact on firm outcomes, we separate all our analysis by whether or not the firm was a competition winner. We are then interested in seeing whether the winners scored more highly by the judges did better than winners with lower scores, and likewise whether the non-winners with higher business plan scores did better than those with lower business plan scores.

Fig. 3 plots the relationship between our four key outcomes (business operation and survival, sales, profits, employment) and the judges' scores for the non-winners. We see that the likelihood of operating a firm three years later is actually slightly lower for those in the upper quintiles of the business plan score. We then fit both unconditional lowess (which codes employment, profits, and sales, as zero for firms not operating) and conditional lowess (which only considers firms in operation) and plot these lines along with the scatterplots for sales, profits and employment. It is apparent that the business plan scores explain very little of the variation in these outcomes, with the fitted lines showing, if anything, poorer results for those with higher scores.

The first column of Table 1 (survival), Table 2 (employment), Table 3 (profits), and Table A1 (sales) then tests whether the relationship between business plan scores and firm outcomes is statistically significant for these non-winners by estimating a logit (survival) or least squares regression (other outcomes) of the outcome on the business plan score

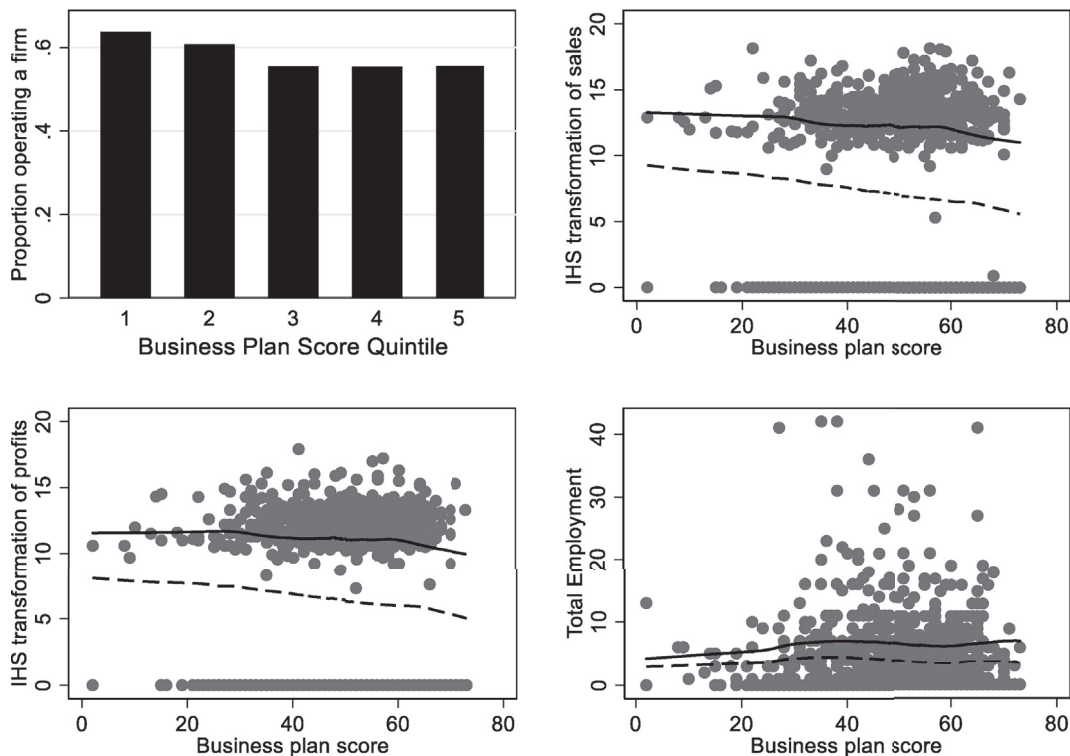


Fig. 3. The Relationship between Business Plan Scores of Judges and Business Outcomes Three Years after applying for Non-Winners in the Business Plan Competition. Notes: Top left panel shows business operation three years after applying by quintile of business plan score. Remaining panels show scatterplots of the sales, profits, and employment three years after applying against the business plan score. The dashed line plotted on each figure is a lowess line of the unconditional relationship, which codes outcomes as zero for applicants who are not operating firms. The solid line on each figure is a lowess line of the relationship conditional on the business being in operation. Employment truncated at 60 workers for readability.

Table 1
Can human experts predict business survival after 3 years?

	Non-Winners			Winners			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Judges	McKenzie	FaWo	Judges	Judges	McKenzie	FaWo
Business Plan score	-0.0002 (0.0014)	-0.0009 (0.0015)	-0.0009 (0.0015)	-0.0000 (0.0006)	-0.0003 (0.0006)	-0.0006 (0.0008)	-0.0006 (0.0008)
Female		-0.1505*** (0.0385)	-0.1272*** (0.0384)			-0.0018 (0.0212)	0.0064 (0.0218)
Age		0.0110*** (0.0032)	0.0085*** (0.0032)			0.0036* (0.0021)	0.0040** (0.0020)
Graduate Education		-0.0211 (0.0632)				0.0262 (0.0362)	
Worked abroad		0.0233 (0.0589)				-0.0447* (0.0243)	
Lottery choice		-0.0194 (0.0292)				-0.0171 (0.0160)	
Wealth		-0.0028 (0.0057)				0.0030 (0.0027)	
Grit		-0.0369 (0.0285)				-0.0031 (0.0154)	
Agriculture		-0.0490 (0.0352)				-0.0156 (0.0196)	
IT		-0.0546 (0.0509)				-0.0264 (0.0244)	
Ability		0.0509*** (0.0169)	0.0333** (0.0162)			-0.0086 (0.0087)	-0.0101 (0.0081)
Manufacture		0.0018 (0.0395)	0.0296 (0.0355)			-0.0259 (0.0203)	-0.0182 (0.0183)
Retail			0.0673 (0.0793)				0.0210 (0.0446)
Self-confidence			0.0083 (0.0091)				0.0062 (0.0050)
Motivation			0.0193* (0.0107)				-0.0054 (0.0048)
Happiness			-0.0031 (0.0120)				-0.0090 (0.0072)
Optimism			0.0332** (0.0136)				0.0072 (0.0082)
Credit			0.0172 (0.0174)				-0.0144** (0.0073)
Existing	0.2125*** (0.0387)	0.1688*** (0.0411)	0.0677 (0.0469)	0.0353** (0.0157)	0.0347** (0.0158)	0.0168 (0.0194)	0.0063 (0.0210)
Log(Total award paid)					0.0248* (0.0145)	0.0219 (0.0140)	0.0183 (0.0154)
Regional Indicators	No	Yes	Yes	No	No	Yes	Yes
Sample Size	1077	1077	1077	1051	1051	1047	1048
P-Value		0.000	0.000			0.530	0.105
Pseudo-R ²	0.024	0.076	0.098	0.011	0.020	0.054	0.064
Accuracy	58.4%	65.6%	66.9%	93.6%	93.6%	93.6%	93.6%

Robust Standard Errors in parentheses. *, **, *** denote significance at the 10, 5, and 1 percent levels respectively. See [Appendix A.1](#) for variable definitions. Dependent variable is having a business in operation 3 years after applying. Marginal effects from logit shown. P-value is for testing that set of controls added by [McKenzie \(2015\)](#) or [Fafchamps and Woodruff \(2017\)](#) (abbreviated as FaWo) are jointly zero.

and an indicator for whether the application was for an existing firm compared to a new firm. The regressions are unconditional, coding outcomes as zero for non-operating firms. For all four outcomes, we see the business score has a small, negative, and not statistically significant association. The R^2 are below 0.04, with even this small amount of variation explained coming largely from the dummy for existing firm status. Judges' scores therefore do not predict which of these non-winners are more likely to succeed.

[Fig. 4](#) then plots the associations between business outcomes and the judges' scores for the winners. We again see the business plan scores explain very little of the variation in these outcomes, with the fitted lines showing a fairly flat relationship for profits and sales, and a slightly upward sloping relationship for employment. Column 4 of [Tables 1–3](#) and [A1](#) then shows the fitted logit and regression relationship, after controlling for whether the application was for an existing firm. The amount of grants the winners received varied among firms, with an average of US\$48,991 and standard deviation of \$17,546. There is a significant positive relationship between business score and amount

received, and more capital has a direct impact on firm outcomes. Therefore, column 5 of [Tables 1–3](#) and [A1](#) also controls for the logarithm of the grant awarded. We see the judges' score has a small and statistically insignificant impact on survival, profits, and sales, with or without this control for grant amount. In contrast, we see in [Table 2](#) that the judges' scores are significantly correlated with employment three years later in these firms, although this relationship weakens and is no longer statistically significant after controlling for the grant amount. Moreover, the share of variation in employment accounted for by these judges' scores is still very low, with an adjusted R^2 of 0.02 in column 4 of [Table 2](#).

This analysis shows that the overall scores of judges are uncorrelated with business success for non-winners, and at most weakly correlated with future employment outcomes, but not other outcomes, for winners. As an additional robustness check, [Table A2](#) shows that this result continues to hold when we condition on judge or region fixed effects.

The overall score aggregates multiple scores on 10 subcomponents, and so it is of interest to examine whether this aggregation masks predictive ability of some subcomponent of the score. In [Appendix Table A3](#)

Table 2
Can human experts predict employment after 3 years?

	Non-Winners			Winners			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Judges	McKenzie	FaWo	Judges	Judges	McKenzie	FaWo
Business Plan score	-0.0023 (0.0168)	0.0008 (0.0182)	0.0046 (0.0180)	0.0881*** (0.0312)	0.0494 (0.0308)	0.0819** (0.0358)	0.0703** (0.0336)
Female		-1.0647** (0.4792)	-0.7637 (0.4865)			-0.6078 (0.8249)	-0.0574 (0.7792)
Age		0.1069*** (0.0413)	0.0775* (0.0398)			0.1705** (0.0823)	0.1541** (0.0775)
Graduate Education		0.9593 (0.9798)				-2.4624*** (0.9219)	
Worked abroad		0.1514 (0.8827)				-1.0069 (1.0370)	
Lottery choice		0.5630 (0.3516)				-0.4718 (0.6622)	
Wealth		0.1012 (0.0685)				0.1315 (0.1338)	
Grit		-0.4658 (0.3613)				-0.4396 (0.6060)	
Agriculture		-0.5865 (0.4563)				-1.0522 (0.7892)	
IT		-1.3288*** (0.4937)				-2.7039*** (0.8412)	
Ability		0.4502** (0.2175)	0.2993 (0.1836)			0.0807 (0.3156)	-0.0850 (0.3168)
Manufacture		-0.4825 (0.4566)	-0.1086 (0.4158)			-0.3916 (0.9174)	0.0146 (0.8037)
Retail			0.6546 (0.9276)				-0.6681 (0.8879)
Self-confidence			0.0726 (0.1164)				0.0990 (0.1854)
Motivation			0.0815 (0.1140)				-0.2382 (0.2048)
Happiness			0.1961 (0.1411)				0.4688* (0.2705)
Optimism			0.1695 (0.1647)				-0.1639 (0.2880)
Credit			-0.0378 (0.1571)				1.0346* (0.5763)
Existing	1.6260*** (0.4510)	1.4832*** (0.4699)	0.3739 (0.5704)	2.5040*** (0.6355)	2.3173*** (0.6226)	2.9320*** (0.7556)	1.7554** (0.8518)
Log(Total award paid)					3.8655*** (0.4328)	3.8208*** (0.4444)	3.8180*** (0.4416)
Regional Indicators	No	Yes	Yes	No	No	Yes	Yes
Sample Size	1077	1077	1077	1051	1051	1051	1051
P-value		0.000	0.000			0.000	0.013
Adjusted R ²	0.012	0.038	0.049	0.020	0.053	0.072	0.079
MSE	34.06	33.18	32.79	103.33	99.77	97.76	97.01

Robust Standard Errors in parentheses. *, **, *** denote significance at the 10, 5, and 1 percent levels respectively. See [Appendix A.1](#) for variable definitions. Dependent variable is number of employees 3 years after applying, coded as zero for applicants without an operating firm. P-value is for testing that set of controls added by [McKenzie \(2015\)](#) or [Fafchamps and Woodruff \(2017\)](#) (abbreviated as FaWo) are jointly zero.

we test whether the ten subcomponents of the business plan score are jointly associated with future outcomes, as well as examining individual subcomponents. For the non-winners, we cannot reject that the ten sub-scores are jointly unrelated to survival, sales, or profits, but we can do so for employment ($p = 0.039$). We also find that the score out of 10 for “ability to manage” is positively and significantly related to all four outcomes - with p-values of 0.003, 0.001, 0.008 and 0.017 respectively - and continues to be so even after applying a Bonferroni or step-down correction for multiple testing. This effect mainly appears to be coming through the extensive margin of predicting whether or not the individual will have a firm in operation. The employment sub-score is also weakly predictive of employment ($p = 0.055$).

For the winners, this “ability to manage” score is not significant for any outcome, and we cannot reject that the ten sub-scores are jointly unrelated to survival, sales, or profits. The employment score does significantly predict employment among the winners ($p = 0.001$), and we can reject that the ten sub-scores are jointly unrelated to the employment outcome ($p = 0.050$). As a reference, a one standard deviation (6.0)

increase in this employment sub-score is associated with the winning firm having 1.1 more workers three years later. Notably, the capstone score - which is intended to reflect the judge’s overall assessment of promise not picked up in the other scores - has no significant association with future business success for either winners or non-winners, and it is even negatively associated with future profits for winners.

3.3. Using human expert judgements to predict business success

An alternative human approach to predicting business success is for economists to use their expert judgement to choose characteristics measured in the baseline survey (or that are measured later but assumed to be time-invariant) that they think may predict outcomes, and then to fit a simple logit or OLS regression model as a function of these characteristics.

This approach was carried out in Appendix 18 in the working paper of [McKenzie \(2015\)](#), but dropped for space reasons in the published version ([McKenzie, 2017](#)). It uses principal components and aggregate index

Table 3
Can human experts predict profits after 3 years?

	Non-Winners			Winners			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Judges	McKenzie	Fafchamps	Judges	Judges	McKenzie	Fafchamps
Business Plan score	-0.0045 (0.0180)	-0.0061 (0.0188)	-0.0066 (0.0183)	-0.0030 (0.0126)	-0.0075 (0.0130)	-0.0093 (0.0143)	-0.0064 (0.0139)
Female		-1.8901*** (0.4955)	-1.5593*** (0.4890)			-1.0142** (0.4454)	-0.8647* (0.4465)
Age		0.1253*** (0.0417)	0.0830** (0.0410)			0.0460 (0.0358)	0.0363 (0.0351)
Graduate Education		-0.9393 (0.8041)				0.0187 (0.6114)	
Worked abroad		0.7856 (0.7839)				-1.3221** (0.6147)	
Lottery choice		-0.0687 (0.3756)				-0.2843 (0.2985)	
Wealth		-0.0222 (0.0742)				0.1071** (0.0461)	
Grit		-0.4456 (0.3614)				-0.1065 (0.2752)	
Agriculture		-0.4053 (0.4597)				-0.2147 (0.3716)	
IT		-0.8472 (0.6560)				-0.4611 (0.4675)	
Ability		0.6695*** (0.2150)	0.4487** (0.2014)			-0.1976 (0.1712)	-0.2346 (0.1570)
Manufacture		0.2671 (0.4931)	0.5600 (0.4426)			0.2026 (0.3869)	0.4204 (0.3590)
Retail			0.9146 (0.9333)				1.5207*** (0.5244)
Self-confidence			0.0781 (0.1148)				0.0782 (0.0921)
Motivation			0.0647 (0.1391)				-0.1942* (0.1035)
Happiness			0.1016 (0.1569)				-0.0817 (0.1253)
Optimism			0.4209** (0.1759)				-0.0094 (0.1376)
Credit			0.2742* (0.1466)				-0.0890 (0.1985)
Existing	2.7393*** (0.4736)	2.4020*** (0.5072)	0.8730 (0.5724)	1.0225*** (0.2961)	1.0015*** (0.2983)	0.7810** (0.3548)	0.5783 (0.4190)
Log(Total award paid)					0.4434 (0.2968)	0.3996 (0.3009)	0.4044 (0.3025)
Regional Indicators	No	Yes	Yes	No	No	Yes	Yes
Sample Size	1058	1058	1058	1036	1036	1036	1036
P-value		0.000	0.000			0.000	0.000
Adjusted R2	0.034	0.081	0.122	0.010	0.011	0.022	0.021
MSE	37.50	35.67	34.08	21.56	21.55	21.31	21.32

Robust Standard Errors in parentheses. *, **, *** denote significance at the 10, 5, and 1 percent levels respectively. See [Appendix A.1](#) for variable definitions. Dependent variable is inverse hyperbolic sine of profits 3 years after applying, coded as zero for applicants without an operating firm. P-value is for testing that set of controls added by [McKenzie \(2015\)](#) or [Fafchamps and Woodruff \(2017\)](#) (abbreviated as FaWo) are jointly zero.

measures for data reduction: a household wealth measure is constructed as the first principal component of 20 durable assets; an ability measure is constructed as the first principal component of the Raven test (itself an aggregate of responses on 12 questions) and digit-span recall scores; and a grit measure is an average of responses on 12 questions proposed by [Duckworth et al. \(2007\)](#). In addition to this, gender, age, attitude towards risk, a dummy for having graduate education, whether the applicant has worked abroad, indicators for the six geographic regions, and indicators for the three most common industries (agricultural business, manufacturing business, IT business) were selected as variables that the literature and an understanding of the business setting suggested might help predict future performance. No model selection criteria were used to select these variables: they were just proposed as an ad hoc judgement of what might matter for predicting business outcomes.

A similar approach is used by [Fafchamps and Woodruff \(2017\)](#) for a business plan competition they ran in Ghana. They propose a set of core measures that they think are likely to be viewed as potentially important: the ability of the owner, past borrowing activity, and management practices, along with controls for gender, age, sector, and location. They

then add attitudinal questions based on the reason for going into business, and on optimism and control over life. Again, no model selection criteria were used to select these variables.

We examine the extent to which these proposed variables coming from baseline survey measures are associated with future business outcomes in the remaining columns of [Tables 1–3 and A1](#). Columns 2 and 6 use the variables selected in [McKenzie \(2015\)](#), while Columns 3 and 7 attempt to match as closely as possible the variables used in [Fafchamps and Woodruff \(2017\)](#). [Appendix A.1](#) describes in detail how these variables were constructed. These regressions allow us to examine whether the proposed models are useful for predicting business outcomes within sample – in the next section we will examine out-of-sample performance to allow comparison with machine learning methods. We also include the business plan score in these regressions so that we can test whether these survey measures add value beyond that contained in the judges' scores, and conversely, whether the judges' scores now contain information once we have conditioned on other features of the applicants.

Consider predicting outcomes for the non-winners. We can strongly reject that the [McKenzie \(2015\)](#) set of regressors, and the [Fafchamps and](#)

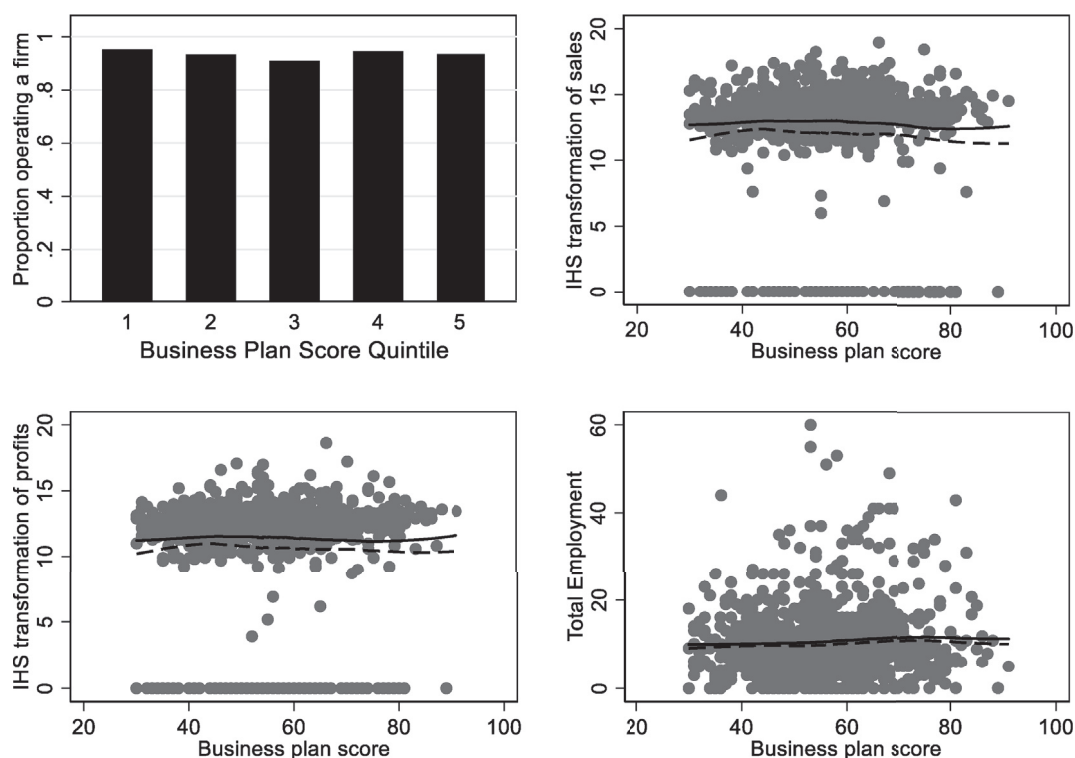


Fig. 4. The Relationship between Business Plan Scores of Judges and Business Outcomes Three Years after applying for Winners in the Business Plan Competition. Notes: Top left panel shows business operation three years after applying by quintile of business plan score. Remaining panels show scatterplots of the sales, profits, and employment three years after applying against the business plan score. The dashed line plotted on each figure is a lowest line of the unconditional relationship, which codes outcomes as zero for applicants who are not operating firms. The solid line on each figure is a lowest line of the relationship conditional on the business being in operation. Employment truncated at 60 workers for readability, profits truncated from below at zero.

Woodruff (2017) set of regressors, are not associated with each of our four business outcomes ($p < 0.001$). This greater performance relative to the judges comes largely from a few characteristics of the owner. First, consistent with a broad body of literature (e.g. Boden and Nucci, 2000; Bruhn, 2009; Robb and Watson, 2012; McKenzie and Paffhausen, 2018), female applicants who do not win are less likely to start or continue businesses than men, and operate firms that are smaller in size and less profitable.

Second, while the business plan competition was designed for youth, it maintained an expansive definition of youth which ranged from age 18 to 40, with a median age of 30 among our non-winners. We see that older applicants were more likely to be running firms three years later among the non-winners, and to run firms of larger size, sales, and profitability. This is consistent with a return to work and life experience, and with evidence from other developing countries that businesses have the lowest failure rates when the owners are middle-aged (McKenzie and Paffhausen, 2018). In contrast to pitching situations where judges see the applicants, the anonymized business plans used for scoring meant that judges did not observe gender or age, and therefore did not consider it in their assessments (and indeed, concerns about sex- or age-discrimination would occur if they did directly consider it).

Third, higher ability of the owner is associated with better performance in all four outcomes for the non-winners. Both the McKenzie (2015) measure, which is the first principal component of Raven test and digit span, and the Fafchamps and Woodruff (2017) measure, which is the first principal component of Raven, digit span, and education, are significant predictors. Examining the sub-components, the Raven test score is the most important contributor here. We then also see some role for business sector in predicting some outcomes: e.g., IT firms tend to hire fewer workers. Our measure of optimism is also positively related with survival, profits, and sales.

These models based on few key variables selected by economists do

better than the judges in predicting outcomes for the non-winners. The business plan scores remain insignificant, even conditioning on the variables in these models. However, these models still only explain a small fraction of the variation in our key outcomes: the adjusted R^2 are 0.04–0.05 for employment, 0.10–0.13 for sales, and 0.08–0.12 for profits (the unadjusted R^2 are also low, always below 0.15).⁴ Moreover, much of the predictive ability comes from the extensive margin of whether firms are operating or not. Appendix Table A4 shows that these R^2 are even lower when we consider only the ability of these predictors to explain the variation in outcomes among firms in operation.

Since almost all of the winners are operating firms after three years (93.5%), there is not much variation in this outcome. We cannot reject that the McKenzie (2015) or Fafchamps and Woodruff (2017) sets of variables are jointly orthogonal to this outcome. We do find some predictive power for employment, sales, and profits, but the R^2 are only 0.07–0.08, 0.02 and 0.02 respectively (the unadjusted R^2 are always below 0.10). Several of the attributes that helped predict outcomes for non-winners are much less predictive of outcomes for winners: gender, while negative, is only significant for profitability and sales; ability is never significant, and has negative coefficients in all but one specification; and being older predicts higher employment, but not higher sales or profits. Sector continues to matter, with IT firms hiring fewer workers, and retail firms having higher profits and sales. We do see that the business plan score of judges continues to be a significant predictor of employment at the 5 percent level, even after controlling for the variables

⁴ By way of comparison, Fafchamps and Woodruff (2017) report an R^2 of 0.14 for profits, 0.23 for sales, and 0.46 for employment. However, their firms had been in operation for an average of 9 years, and do not include start-ups. Moreover, they are able to include baseline sales and profits among their predictors, which exhibit some persistence and which we do not have for start-up firms.

in these models.

It is also notable from Tables 1–3 and A1 that some characteristics that were expected to help distinguish among business owners did not have significant predictive power. A first example is education, which is negatively associated with employment among winners, and otherwise not significant. One possible reason for not seeing the expected positive return to education (e.g. Michelacci and Schivardi, 2016; Queiro, 2016) is that the earlier stages of the competition (including the need to apply online) appear to have selected heavily on education. Only 5.5% of Nigerian youth have university education, but 52% of the applicants, and 64% of those invited to business plan training had this level. Likewise, applicants who got through to the business plan submission stage are wealthier, and more likely to have overseas work experience, than the average Nigerian youth, yet conditional on getting to the business plan submission stage, we generally do not find these variables to be significant predictors of business outcomes.⁵ The literature has also suggested risk preferences determine selection into entrepreneurship and the success of these businesses (e.g. van Praag and Cramer, 2001; Hvide and Panos, 2014; Skriabikova et al., 2014), yet we see no significant impact of our risk preference measure. Nor do we see any predictive power for the grit measure of Duckworth et al. (2007).

4. Machine learning approaches to predicting business success

The human prediction approaches rely on a small set of variables, and on methods like principal components and simple aggregation for data reduction. Yet, in practice, the application and baseline data contain a much richer set of possible predictors. Many of these come from just considering a lot more variables, but the number of predictors becomes even larger once we consider how responses to certain questions should be coded. We discuss a few specific examples, and then the total number of potential predictors we consider.

A first example comes from questions with Likert scales. For example, applicants were asked a set of self-efficacy questions about their confidence in carrying out nine different business tasks, such as “find and hire good employees to expand your business”. They had six possible responses to their confidence on this task: not at all confident, somewhat confident, quite confident, very confident, don’t know, or not applicable. Rather than trying to code these as binary variables and then aggregate, an alternative is to code each task and response as a separate dummy variable (e.g. being not at all confident that they can estimate accurately the costs of a new project). This yields 54 potential predictors.

A second example comes from the Raven test. Applicants were given 12 different puzzles to solve, and had a choice from among 8 options each time, with one of these options correct. The standard approach, used in the models estimated in the previous section, is to aggregate these into a single score out of 12. Yet some questions may be potentially more useful in predicting business success than others, and different types of wrong answers may reflect different abstract reasoning deficiencies. For example, question 6 had 39% choose the correct answer, five other answers were chosen by 9%–11% of individuals each, and the remaining two answers were all chosen by between 5% and 7% each. Coding each question and answer as a separate dummy variable then yields 96 potential predictors.

A third example comes from the business sector that the applicant proposes for their new business venture, or has for their existing business. They were asked to select from 32 different business sector categories. In the ad hoc model of McKenzie (2015), four of these were aggregated together to form a manufacturing dummy, and two other sectors were

⁵ Wealthier individuals have higher profits and sales as winners (at the 5 percent significance level), while we find overseas work experience has a negative association with profits and survival for winners (at the 10 percent level), while these variables are not significant for other outcomes for winners, or for any outcome for non-winners.

also chosen, while the rest were lumped together as the base comparison group. In contrast, we consider all 32 categories as possible predictors here.

However, the increase in the number of variables considered does not just come from dis-aggregation of variables into more categories, but also from considering a wide range of additional variables collected during the initial application and business plan submission. As noted in Section 2.2 and detailed in Appendix 1, these include measures of answer quality, family composition, languages spoken, how the business was named (Belenzon et al., 2017; Guzman and Stern, 2017), business and employment background of the owner, registration and tax status of existing firms, and the owner’s projections of firm growth. These are all variables which may plausibly be correlated with business success.

Following Mullainathan and Spiess (2017), we do not drop redundant variables, e.g. aggregate indexes such as the grit or ability, since they could be useful to obtain better predictions with less complexity. The net result of this is that there are 566 possible predictors that we can consider. This figure includes the initial application score and the ten different subcomponent scores from the judges, in addition to the survey data. In order to make a fair comparison, we initially do not include the business plan scores as inputs in the machine learning algorithms. Moreover, if we also start considering interaction terms between some of these predictors, the number of variables to consider can easily exceed the number of firms for which we have data. It is then not possible to include all of these variables in a standard OLS regression, or logit model. Machine learning provides tools for using these high-dimensional data.

4.1. The machine learning algorithms

We use three different machine learning approaches. We briefly summarize each here, and provide more details in Appendix A.3. A more comprehensive review of the tools available to practitioners is provided by Hastie et al. (2009), Ng (2016), as well as Mullainathan and Spiess (2017).

The first method is **LASSO** (Least Absolute Shrinkage and Selection Operator). This method adds a penalization term to the OLS objective function:

$$\hat{\beta}(\lambda) = \operatorname{argmin}_{\beta \in \mathbb{R}^k} \sum_{i=1}^n (y_i - x_i' \beta)^2 + \lambda \sum_{j=1}^k |\beta_j|$$

where k is the number of potential predictors (potentially larger than the number of observations, n), and λ is the penalization term. In practice, this selects variables with high predictive power, and shrinks their coefficients, while constraining all other variables to have coefficients of zero. As explained in Section 4.2, we choose the penalization term λ via cross-validation.

The second method is **Support Vector Machines (SVM)** for our business operation outcome, and the extension of **Support Vector Regression (SVR)** for our continuous outcomes of employment, sales, and profits (Guenther and Schonlau, 2016). SVM aims to classify the data into two groups (in our case, into firms that will be operating and those that will not) by means of a hyperplane in a high-dimensional space. It can be written as a modified penalized logistic regression that solves the following objective function:

$$\hat{\beta}(C) = \operatorname{argmin}_{\beta \in \mathbb{R}^k} C_1 \left[\sum_{i=1}^n y_i \max\{0, 1 - K_i' \beta\} + (1 - y_i) \max\{0, K_i' \beta - 1\} \right] + \beta_2$$

There are four terms here that differ from that in the standard logistic regression, and which allow SVM greater flexibility in covering a wide range of functional forms and dealing with high-dimensional data. In addition to the penalization factors, C_1 and β_2 ; there is a kernel K which controls the smoothness of the fitted curve, and a “soft margin” that comes from the $\max(0, \cdot)$ component, which allows for some observations

to be misclassified close to the separating margin. We consider a Gaussian kernel, and use cross-validation to choose the penalization term and the kernel smoothing parameter.

SVM has been found to achieve better performance than a number of other machine learning algorithms in some applications (Maroco et al., 2011). Nevertheless, given the similarity in terms of objective function, in other cases performances have been found to be similar to those obtained with logistic regressions (Verplancke et al., 2008). One downside of SVM and SVR are that the parameters are difficult to interpret, and so we will not be able to say which variables matter most for determining the predictions made by this method.

The final method we consider is **Boosting** (also called Boosted Regression), using the gradient boosting algorithm of Friedman et al. (2000). Gradient boosting is an ensemble method, which builds a classifier out of a large number of smaller classifiers. Like random forests, it works by combining many regression trees together. Key differences are that boosting does this sequentially, where in each iteration, observations that were misclassified by the previous classifier are given larger weights; and that boosting typically uses shorter trees, including tree stumps (a single split of the data), aggregating over more of them than random forests.⁶ We consider trees with up to 8 splits, thus allowing up to 8-way interactions between variables in a very flexible way.

Performance is improved by using bagging, which uses only a random subset of the training data set to build the tree in each iteration. To avoid over-fitting, i.e. the risk of estimating a model which fits a training sample very well, but may fail to properly work with new samples, it is also possible to introduce a shrinkage parameter. This reduces the contribution of each additional tree, thus decreasing the impact of an over-fitted tree. Boosting has been found to have superior performances than a number of other machine learning algorithms in many simulations (Bauer et al., 1999; Friedman et al., 2000) and was used by Chalfin et al. (2016) in their work on predicting police hiring.

4.2. Estimation by machine learning

We follow the recommendation of Mullainathan and Spiess (2017) and split the data into two sub-samples (separately for winners and non-winners). A training sample (80% of the data) is used to calibrate and estimate the algorithm under each of the three methods. Out-of-sample performance is reported using the hold-out sample (the remaining 20% of the data).

The algorithms are calibrated using 5-fold cross-validation.

1. Divide the 80% training sample in 5 folds.
 - a. Select a possible numeric value for each parameter (e.g. 0.08 for λ in LASSO).
 - i. Train the algorithm on 4 of the 5 folds using the selected parameter values.
 - ii. Predict the outcome variable (e.g. profits) for the firms in the remaining fold and compute the relevant statistics (e.g. mean squared error).
 - iii. Repeat the above procedure five times, one for each fold.
 - iv. Compute the average performance in the 5 folds.
 - b. Repeat for each possible combination of the parameter values.
 - c. Select the combination of the parameter values that minimizes a given loss function (e.g. minimizing the mean squared error). These in-sample statistics are reported in Appendix A.4.
2. Train the algorithm using the 80% training sample with the selected parameter values.
3. Predict the outcome variable for the firms in the 20% hold-out sample and compute the relevant statistics (e.g. out-of-sample mean squared

⁶ Gradient boosting appears to perform slightly better than random forests in dimensions of 4000 or fewer predictors, but requires more tuning (Zygmunt, 2016).

Table 4
Prediction (out-of-sample) accuracy for business survival among non-winners.

	Model	Predictors	Accuracy		
			Mean	C.I.	
Judges	1	Logit	Constant	58.8%	[52.3%, 65.3%]
	2	Logit	Application Score	58.8%	[52.3%, 65.3%]
	3	Logit	Business Score	58.8%	[52.3%, 65.3%]
	4	Logit	SubScore	56.9%	[50.5%, 63.4%]
Single Predictor	5	Logit	Gender	58.8%	[52.3%, 65.3%]
	6	Logit	Age	63.0%	[56.6%, 69.4%]
	7	Logit	Necessity Firm	55.1%	[48.3%, 61.9%]
	8	Logit	Grit	61.6%	[55.2%, 67.9%]
	9	Logit	Digit-span Recall	60.6%	[54.3%, 67.0%]
	10	Logit	Raven Test	67.1%	[61.0%, 73.2%]
	11	Logit	Registration time	58.8%	[52.3%, 65.3%]
Economist Models	12	Logit	McKenzie	63.9%	[57.4%, 70.4%]
	13	Logit	FaWo	67.1%	[60.7%, 73.6%]
	14	Logit	FaWo + BusScores	66.2%	[59.8%, 72.6%]
	15	OLS	FaWo + BusScores	65.7%	[59.4%, 72.1%]
Machine Learning	16	Probit	FaWo + BusScores	65.7%	[59.4%, 72.1%]
	17	LASSO	All	60.2%	[53.7%, 66.7%]
	18	SVM	All	66.2%	[59.8%, 72.6%]
	19	Boosting	All	63.9%	[57.6%, 70.2%]

Notes: outcome is whether an applicant who did not win the business plan competition is found to be operating a firm three years later. Out-of-sample accuracy is the ratio of true positives and true negatives to all observations, and is computed using the 20% hold-out sample (N = 216). The numbers in brackets are bootstrapped 95 percent confidence intervals for hold-out prediction performance. Judges use scores from competition judges to predict the outcome. Single predictor models use a single survey measure. The economist models are the models of McKenzie (2015) and Fafchamps and Woodruff (2017) (abbreviated as FaWo) used in Tables 1–3 All models apart from the constant model also include a dummy for whether the applicant was for an existing firm compared to a new firm.

error). These are the out-of-sample statistics reported in Tables 4–6 and A10.

4. Compute 95% confidence intervals for these statistics using bootstrapping. It is important to mention that these are not the standard confidence intervals computed in regression models. As emphasized in Mullainathan and Spiess (2017), “these uncertainty estimates represent only variation of the hold-out sample for this fixed set of prediction functions, and not the variation of the functions themselves”.

For the algorithms that do not require any tuning, such as OLS and Logit, we simply follow steps 2–4. Since almost all winners were operating firms, we only predict operation and survival for non-winners, giving seven outcomes to predict.

4.3. Assessing accuracy

We follow standard practice in using the mean squared error (MSE) as the main criteria to compare the accuracy of different models for our continuous outcomes (employment, sales, and profits). For the sake of completeness, we have also reported the square of the Pearson correlation coefficient, i.e. the correlation between the actual and fitted dependent variable. This is equivalent to the R^2 in linear least squares regressions.

In contrast, there are several different measures used in the literature when considering the goodness of fit for binary outcomes, such as business operation and survival in our case. The starting point is usually the matrix comparing predicted outcomes with actual ones:

		Predicted values	
		0	1
Actual values	0	True Negative (TN)	False Positive (FP)
	1	False Negative (FN)	True Positive (TP)

Table 5
Prediction (out-of-sample) accuracy for total employment.

	Model	Predictors	Non-winners			Winners			
			MSE		R ²	MSE		R ²	
			Mean	C.I.		Mean	C.I.		
Judges	1	OLS	Constant	1.71	[1.54, 1.87]	0.0%	0.94	[0.68, 1.20]	0.0%
	2	OLS	Application Score	1.68	[1.50, 1.86]	1.9%	0.86	[0.61, 1.10]	8.9%
	3	OLS	Business Score	1.68	[1.50, 1.86]	1.9%	0.85	[0.61, 1.10]	9.1%
	4	OLS	SubScore	1.65	[1.48, 1.83]	3.2%	0.86	[0.61, 1.11]	8.2%
Single Predictor	5	OLS	Gender	1.65	[1.47, 1.83]	3.2%	0.92	[0.66, 1.18]	1.9%
	6	OLS	Age	1.64	[1.46, 1.82]	3.9%	0.91	[0.66, 1.16]	3.3%
	7	OLS	Necessity Firm	1.70	[1.52, 1.88]	1.3%	0.92	[0.67, 1.18]	1.5%
	8	OLS	Grit	1.67	[1.48, 1.86]	2.4%	0.93	[0.67, 1.19]	1.3%
	9	OLS	Digit-span Recall	1.69	[1.50, 1.88]	1.7%	0.93	[0.67, 1.18]	1.5%
	10	OLS	Raven Test	1.61	[1.43, 1.80]	5.6%	0.94	[0.68, 1.20]	0.6%
	11	OLS	Registration time	1.68	[1.50, 1.85]	2.0%	0.93	[0.67, 1.20]	0.8%
Economist Models	12	OLS	McKenzie	1.59	[1.40, 1.78]	6.9%	0.91	[0.66, 1.17]	2.9%
	13	OLS	FaWo	1.54	[1.35, 1.72]	10.3%	0.93	[0.66, 1.20]	2.2%
	14	OLS	FaWo + BusScores	1.54	[1.35, 1.72]	10.3%	0.86	[0.60, 1.13]	8.1%
	15	OLS	FaWo + Baseline	1.56	[1.38, 1.74]	9.0%	0.86	[0.61, 1.11]	8.3%
Machine Learning	16	LASSO	All	1.54	[1.38, 1.70]	10.6%	0.85	[0.60, 1.09]	15.1%
	17	SVM	All	1.57	[1.39, 1.76]	8.4%	0.88	[0.64, 1.12]	6.2%
	18	Boosting	All	1.53	[1.35, 1.71]	10.3%	0.91	[0.65, 1.16]	3.8%

Notes: outcome is the inverse hyperbolic sine of total employment in the firm three years after applying, coded as zero for applicants not operating firms. Models are estimating separately for competition non-winners and winners. MSE is the out-of-sample mean-squared error computed using the 20% hold-out sample. The numbers in brackets are bootstrapped 95 percent confidence intervals for hold-out prediction performance. R² is the squared correlation between the fitted outcome and actual outcome in the 20% hold-out sample. Judges models use scores from competition judges to predict the outcome. Single predictor models use a single survey measure. The economist models are the models of McKenzie (2015) and Fafchamps and Woodruff (2017) (abbreviated as FaWo) used in Tables 1–3 All models apart from the constant model include a dummy for whether the applicant was for an existing firm compared to a new firm. Models 2–4 and 14 for winners also control for the total award received (in logs).

Table 6
Prediction (out-of-sample) accuracy for monthly profits.

	Model	Predictors	Non-winners			Winners			
			MSE		R ²	MSE		R ²	
			Mean	C.I.		Mean	C.I.		
Judges	1	OLS	Constant	38.89	[37.18, 40.61]	0.0%	20.38	[15.37, 25.39]	0.0%
	2	OLS	Application Score	39.02	[36.41, 41.62]	1.0%	20.32	[15.36, 25.28]	0.6%
	3	OLS	Business Score	39.04	[36.41, 41.67]	1.0%	20.30	[15.35, 25.24]	0.7%
	4	OLS	SubScore	38.90	[36.14, 41.65]	1.3%	20.18	[15.38, 24.99]	1.3%
Single Predictor	5	OLS	Gender	38.00	[35.34, 40.67]	2.6%	20.39	[15.45, 25.33]	0.5%
	6	OLS	Age	38.71	[35.98, 41.44]	1.5%	20.04	[15.16, 24.91]	1.8%
	7	OLS	Necessity Firm	39.26	[36.62, 41.90]	0.8%	19.92	[15.11, 24.74]	2.7%
	8	OLS	Grit	38.91	[36.25, 41.56]	1.2%	19.88	[15.06, 24.71]	2.8%
	9	OLS	Digit-span Recall	38.57	[35.90, 41.24]	1.7%	19.80	[14.98, 24.61]	3.4%
	10	OLS	Raven Test	37.28	[34.37, 40.19]	4.4%	19.89	[15.04, 24.75]	2.5%
	11	OLS	Registration time	39.15	[36.49, 41.80]	0.9%	19.98	[15.12, 24.83]	2.1%
Economist Models	12	OLS	McKenzie	37.99	[34.51, 41.48]	4.0%	20.90	[15.83, 25.98]	0.4%
	13	OLS	FaWo	34.54	[30.75, 38.33]	11.2%	20.79	[15.80, 25.77]	0.6%
	14	OLS	FaWo + BusScores	34.55	[30.78, 38.31]	11.2%	21.04	[16.00, 26.08]	0.3%
	15	LASSO	All	35.23	[32.86, 37.60]	11.1%	20.32	[15.30, 25.34]	0.9%
Machine Learning	16	SVM	All	33.79	[30.50, 37.08]	13.1%	20.07	[14.85, 25.29]	1.3%
	17	Boosting	All	35.41	[31.51, 39.32]	9.2%	21.07	[16.05, 26.08]	0.1%

Notes: outcome is the inverse hyperbolic sine of monthly profits in the firm three years after applying, coded as zero for applicants not operating firms. Models are estimating separately for competition non-winners and winners. MSE is the out-of-sample mean-squared error computed using the 20% hold-out sample. The numbers in brackets are bootstrapped 95 percent confidence intervals for hold-out prediction performance. R² is the squared correlation between the fitted outcome and actual outcome in the 20% hold-out sample. Judges models use scores from competition judges to predict the outcome. Single predictor models use a single survey measure. The economist models are the models of McKenzie (2015) and Fafchamps and Woodruff (2017) (abbreviated as FaWo) used in Tables 1–3 All models apart from the constant model include a dummy for whether the applicant was for an existing firm compared to a new firm. Models 2–4 and 14 for winners also control for the total award received (in logs).

An overall measure of the goodness of fit is the accuracy rate, which is the proportion of predictions that are correct out of all observations:

$$Accuracy = \frac{TP + TN}{Total\ number\ of\ observations}$$

The recall rate is also commonly used in the case of imbalanced data, i.e. when the number of positive values is much smaller or larger than the number of negative values:

$$Recall\ (or\ Sensitivity) = \frac{TP}{TP + FN}$$

Both LASSO and Boosted regression predict probabilities for binary outcomes (in contrast, SVM classifies observations as 1s or 0s, and does not produce probabilities). We follow convention of predicting an outcome of one (survival) when the predicted probability is greater than or equal to 0.5, and zero otherwise. This is in line with the Bayes classifier (Hastie et al., 2009): the accuracy rate is maximized by assigning each

observation to the most likely class, given its predicted probabilities.

Subrahmanian and Kumar (2017) and Sansone (2018) note that the choice of criteria should depend on context and whether there are strong reasons to prefer the ability to detect true positives versus to avoid false negatives. In the context of predicting firm operation, we do not believe there are strong reasons to prefer one versus the other, and so use accuracy as our measure of performance. In Section 5.3, we consider the ability of these methods to predict the upper tail of performance on profitability, in which case we also report the recall rate.

To compare the prediction success of our machine learning approaches to those of the judges and of the models of economists shown in Section 3, we also use the same 80% training samples to fit these models, and then the 20% hold-out sample to get predictions, MSE, and accuracy rates. We fit the models shown in Tables 1–3, which predict outcomes as a function of the judges score, the McKenzie (2015) predictors, or the Fafchamps and Woodruff (2017) predictors. Unless otherwise noticed, we do not include the judge score when fitting the McKenzie (2015) or the Fafchamps and Woodruff (2017) models. For robustness, we consider two other models based on the judges scores: using just the first-round application score, or using all ten sub-scores (as in Appendix Table A3). We show the performance of the Fafchamps and Woodruff (2017) predictors with and without the business plan score included.

Finally, we also compare to the performance of seven models each based on just a single predictor (gender, age, necessity firm, grit, digit-span, Raven test, registration time for the competition⁷). It is worth emphasizing that these seven variables have not been selected at random, but based on previous research citing them as important predictors of future success. Therefore, these univariate models can be seen as simpler version of the models developed by McKenzie (2015) or Fafchamps and Woodruff (2017).

4.4. Results from machine learning and human predictions

Table 4 examines the out-of-sample performance of the different models in predicting business operation and survival for the non-winners three years after they applied to the competition. As a benchmark, the first row reports the out-of-sample accuracy rate (58.8%) of a model which contains only a constant. The next three models consider the performance of the scores from judges – the initial application score (model 2), the business plan score (model 3), or the ten sub-scores from the business plan (model 4). All three are no more accurate than just using the constant model, re-iterating the poor performance we saw in Table 1.

Models 5 through 11 just use a single predictor, and most of these models perform better than the judges. The best performance here comes from the Raven test, which is able to correctly predict whether or not a business will be operating in 67.1% of the cases. Models 12 through 16 consider the ad hoc models of the human experts described in Section 3.3. Both the McKenzie (2015) and the Fafchamps and Woodruff (2017) models have relatively high performances, with 63.9% and 67.1% accuracy respectively. We see that the accuracy of the Fafchamps and Woodruff (2017) does not improve when we also add the business plan score, and is similar whether we use probit model or a linear probability model instead of a logit function. The Fafchamps and Woodruff (2017) model provides a 14% improvement in accuracy compared to only using the business plan score of the judges, and its lower bound in the 95% confidence interval is above the mean accuracy obtained using the judges' scores. However, the Fafchamps and Woodruff (2017) model does no better than just using the Raven test score alone.

Finally, models 17 through 19 are the machine learning models. Their out-of-sample accuracy rates range from 60.2% for LASSO to 66.2% for

SVM. In-sample 5-fold average accuracy rates are close to these out-of-sample accuracy rates, thus reducing any concern about over-fitting (Table A11). The ranges used to tune the ML parameters are appropriate for the sample used: selected values of the model parameters (such as λ in LASSO) are in the interior of the intervals used during cross-validation (Tables A6–A9). Fig. A2 further shows that the cross-validation procedure identifies the best λ in LASSO, and that the choice of such parameter deeply affect the in-sample performances of the algorithm. Therefore, even if properly calibrated, these high-dimensional modeling approaches do no better than the ad hoc models of human experts, or than just using the Raven test score alone, in terms of predicting business survival among the non-winners. Moreover, none among these 19 models achieves outstanding performances: the upper bounds in the 95% confidence intervals do not exceed a 73.6% accuracy rate.

Tables 5–6 and A10 then shows the out-of-sample MSE and square of the Pearson correlation coefficient for employment, profits, and sales, for both the non-winners and winners. Consider the non-winners. For both employment (Table 5) and sales (Table A10), Boosting, LASSO and the Fafchamps and Woodruff (2017) model have the lowest MSE, followed by SVR. The Raven test alone is also a powerful predictor. The business plan scores of judges are among the worse. Nevertheless, the differences are small across models, with only a 9% reduction in MSE for employment obtained when moving from using the judge's business plan scores to the lowest MSE model (that of Boosting). SVR and the Fafchamps and Woodruff (2017) models have also the lowest MSE for profits. The squared correlations are below 11% for employment, while they are 13% or below for sales and profits.

Among the winners, we do not see large differences in performance across the different methods, especially for sales and profits. The MSE of the judges, human expert models, and machine learning models are within 7% of one another for these two outcomes for winners, and the squared correlations are always below 4%. We see more differences in performance for the employment outcome. The best performance comes from LASSO, with a squared correlation of 15%. For this outcome, the judge scores have superior performances than the other parsimonious models.

As with business survival, the Online Appendix includes several additional results to confirm that the ML algorithms are properly calibrated. Tables A6–A9 show that the selected model parameters in LASSO, SVM and Boosting are in the interior of the intervals used during cross-validation. Figs. A3–A4 provide further examples of how cross-validation selects the optimal value of λ in LASSO, and that different values lead to large differences in MSE. In-sample performances are reported in Tables A12–A14.

4.5. Does machine learning use different predictors from human experts?

The machine learning models potentially select different predictors in each fold, and for each outcome. Nevertheless, it is of interest to see whether the machine learning algorithm ends up putting high weight on variables that would not typically be considered by experts when predicting business success. SVM and SVR use kernels over many variables and so are not easily interpretable in this regard.

Friedman (2001) and Schonlau (2005) note that for boosting regression, one can compute the influence of a variable. This depends on the number of times a variable is chosen across all iterations (trees) and its overall contribution to the log-likelihood function. Such values are then standardized to sum up to 100. We therefore look at the variables which have been selected at least once in the 6 status-outcome estimations (non-winners/winners, employment/sales/profits) reported in Tables 5–6 and A10. It is interesting to note that, among the 566 predictors considered, 182 have been picked by the algorithm to construct a tree. However, around 95 of them have been selected only once, 87% of them have been selected 3 times or less, and only 24 variables have been used at least 4 times.

Table A20 lists these 24 predictors along with the number of boosted

⁷ Banerjee and Duflo (2014) suggest enrolling late is a marker for being disorganized and less able to complete tasks on time, which we hypothesize could be bad for business growth.

regression they have been used in, and the influence for each status-outcome. The most striking result is that most of these regressors are similar to the ones selected by heuristic models. Indeed, the list includes demographic characteristics, wealth and income indicators, measures of self-confidence, optimism and happiness, number of employees at baseline, Raven test, grit and overall ability. The most frequently selected variable is the ability measure used in [Fafchamps and Woodruff \(2017\)](#), followed by registration time. Given these results, it is perhaps not surprising that the machine learning algorithms do not produce substantial improvements over the other models considered: the inputs selected by the models end up being very similar, and most of the boosted regressions select only 1 or 2 splits, implying low levels of interactions among these variables. In other words, the simple models of human experts already contain many of the most powerful predictors, and there seems not to be important nonlinearities in this context.

Nevertheless, the list also includes some additional variables which may be useful to improve predictions. Taxes paid in the past and future expectations about the amount of money needed to expand are often selected. Experience abroad seems to be relevant as well. Family links and household composition can also play a role: the list includes the number of brothers and sisters. The length of the first answer is often chosen. Finally, one specific answers on the Grit scale is chosen.

The same exercises can be replicated for LASSO. Among the 566 possible predictors, only 51 variables were selected at least once by the LASSO algorithms reported in [Tables 5–6](#) and [A10](#), but 19 of them were selected two or three times. [Table A21](#) lists these predictors. There are some overlaps with the variables selected in [Table A20](#), e.g. age, gender, and some indicators of self-confidence, optimism, ability and family composition. Quite interestingly, one incorrect answer from the Raven test is selected, thus suggesting that there might be important information in not just whether individuals get the right answer, but also in which wrong answer they choose.

It is worth emphasizing that, as discussed in [Mullainathan and Spiess \(2017\)](#), that different algorithms and different samples may lead to different variable selections. Indeed, if some variables are highly correlated, then they be used as substitutes when predicting business success. The final set of selected variables depends on the specific finite sample used to train an algorithm. Nevertheless, it is remarkable that several variables are selected by both LASSO and Boosting. This supports the conclusion regarding their high predictive power. Moreover, the aim of this section is to identify top predictors, and to test whether this set of variables overlaps with those selected by economists. We are not making any causality claim.

We investigate two further questions regarding the comparison of machine learning to other approaches. The first is whether the judges reading of the business plans and subsequent scoring do provide valuable predictive information that can not otherwise be picked up by the machines. [Table A16](#) shows that, except when predicting survival, business plan scores and sub-scores are not selected at all by LASSO when added to the set of possible predictors. This suggests that any information in the business plan scores is either predictable by other variables used in the model, or else of little predictive value. The second question is whether any differences in performance between the ad hoc models of economists and the machine learning variables reflect differences in the variable inputs, or the ability of machine learning methods to combine these with more flexible functional forms. [Table A17](#) compares the performance of the [Fafchamps and Woodruff \(2017\)](#) model to our full SVM and Boosting models, and to SVM and Boosting models that only start with the same input variables as Fafchamps and Woodruff, with the differences then coming from functional form flexibility these models provide. We see the performances are very similar across models, but where small differences arise (e.g. sales and employment for winners), they appear to stem more from the machine learning models using different inputs than from their use of more flexible functional forms.

5. Why can't we predict better, and how important is this economically?

Comparing all of these sets of results leads to several conclusions. First and foremost, predicting outcomes for these businesses is difficult, with the out-of-sample performance of all models quite low, especially for the profits and sales of winners. Second, machine learning does not consistently lead to improved performance compared to human predictions in this context. In fact, simple models from human experts slightly outperform at least one out of three machine methods for every outcome considered. Confidence intervals are also wide and they often overlap across different models. However, some of the machine learning predictions do yield on average more accurate predictions than using the judge scores, and are often among the best performing models. Third, a simple measure of cognitive skills, the Raven test, performs at least as well on average as using judge scores, and has often similar performances to those of human expert models, or machine learning. Fourth, the scores from the initial applications almost always provide a lower MSE than the business plan scores, raising doubts about the cost-effectiveness of this added time and cost.

These results raise questions of why the business plan scores from judges were not more powerful in predicting success, and of why the machine learning did not perform better. We examine each in turn, and then consider whether the methods do better at the very top tail.

5.1. Why didn't the human experts do better?

While there are several plausible explanations as to why the judges' scores are not more predictive of business success, our data enables us to rule out a number of these.

A first potential explanation is that perhaps the judges scores do not reflect which firms the judges thought would succeed most for one of two possible reasons. A first reason could be if individual judges differed in how important they viewed the different components of the overall score in determining business success, but could not reflect this because they had no control over the weights used. However, as [Table A3](#) shows, most of the sub-scores also do not predict performance, and, in particular, the capstone score, which best reflects their overall assessment of whether there is something special about the firm, is also not a good predictor of performance. A second reason could be if judges had instead wanted to choose firms on the basis of which would benefit most from the grants (who had the largest treatment effects), instead of which firms would be most successful. However, [Appendix 18 of McKenzie \(2017\)](#) finds that, if anything, new firm applicants scored more highly by judges had lower treatment effects over the first year, and by the time of the three-year horizon that we use here, there is no significant treatment heterogeneity according to the business plan score.

A second possible explanation for the poor performance of the judges could stem from them only observing anonymized plans, and so not being able to statistically discriminate on the basis of characteristics like gender and age that our analysis has found to have predictive power. To examine this possibility, in [Table A15](#) we drop age and gender from the economist models and Lasso models, and find that these models still do better than the judges in predicting out-of-sample performance, even without this statistical discrimination.⁸

A third possibility is that, unlike the economist models and machine learning algorithms, the judges did not have access to the training dataset with outcomes for a subset of applicants. However, the judges did have their past experience of the Nigerian business environment and working

⁸ All models continue to perform poorly when predicting business success among winning firms, especially for profits and sales. It would also be an interesting exercise in future research to test whether judges would change how they scored proposals if age and gender information were added to the plans, as it is unclear whether they would statistically discriminate.

with similar firms, and the lessons of their fellow judges discussed in judge training and codified in the scoring rubric. Moreover, even if they had observed outcomes, it may have been difficult for them to understand what factors mattered most for business success. As an example, McKenzie (2018) shows that even after having gone through the YouWin program, firm owners were unable to tell how much winning the program had mattered to their firm outcomes.

Fourth, perhaps the problem is that it is much easier for judges to distinguish a bad business plan from a good one, than it is for them to distinguish among top business plans. Our data only comes from those who go through to the stage of submitting a business plan, who are the top-25 percent of applicants. But even among this group, the judges determined some of the business plans to be particularly poor, giving scores as low as 3 out of 100. If judges found it easier to distinguish the bottom firms from the rest, we should expect to see non-linearities in the relationship between score and outcomes, with at least a strong positive relationship among lower-scored proposals. Fig. 3 shows this is not the case.

5.2. Why didn't machine learning do better?

Likewise, there are a number of potential explanations as to why the machine learning did not have better predictive performance, and we can use our data to explore the plausibility of each of these. These concerns relate to the sample size, inputs into the models, and implementation of the machine learning methods.

A first potential explanation could be that the machine learning needs a larger sample size to perform well. Although we cannot see whether increasing the sample would improve performance, we can examine how sensitive our results are to decreasing the training sample size used. We start with our full 80 percent training sample (recall this is just over 800 observations), and then randomly remove a fraction of this training sample to see how stable the out-of-sample accuracy rate is. The results are shown in Fig. A5. Prediction accuracy for non-winners is substantially worse when using only 30 percent or less of the training sample (samples of under 250 observations), continues to improve as we add more data, and is then reasonably stable when using only 80 or 90 percent of our training sample (samples of 600 or 700 instead of 800): the out-of-sample MSE is 1.55 when using 80 percent of the training sample, 1.54 when using 90 percent, and it remains 1.54 when using the full training sample to calibrate our LASSO algorithm. Similar conclusions can be reached when repeating the same exercise for winners (Fig. A6): if anything, the out-of-sample MSE increases when moving from using 90 percent of the training sample to 100 percent (from 0.837 to 0.847). This suggests that small increases in sample size would not improve substantially the model performance, but of course we cannot rule out that performance would be better if data on millions of firms were available. Nevertheless, we note that our sample of over 2000 firms is substantially larger than is possible for most business plan competitions, or than the information sets of most venture capitalists looking to make investments.

A second set of concerns are based on the inputs into the models. A first issue might be that the inputs into the model may lack predictive power, if respondents strategically mis-report items on their application in order to enhance their perceived chances of winning the competition. To the extent that this is occurring, we view it as a feature, not a bug, since we are using data collected under competition conditions to predict which competition applicants will succeed.⁹ Nevertheless, we do not have strong reasons to expect strategic reporting to greatly affect our machine-learning results, in part because it is unclear for many of our variables in which direction a strategic applicant should report (e.g. is it better to report having overseas work experience or not? is it better to say

you are risk-seeking or not? Is it better to say your household has a freezer or not?).

Another issue concerning the inputs is the possibility that the large number of additional input variables we consider for machine learning are highly correlated with the simple set of variables used in our ad hoc economist models, and so do not contain useful independent information. To investigate this possibility, we regress each machine learning input on the McKenzie (2015) or Fafchamps and Woodruff (2017) variable sets, using an OLS for non-binary variables and a Probit for binary variables. Fig. A1 shows kernel densities of the resulting R^2 (or pseudo- R^2 for binary variables), for winners and non-winners separately. While there is a mass close to $R^2 = 1$ (there is some overlapping between the two sets), most of the R^2 are below 0.2, showing that there is indeed substantial independent information beyond the set of variables used in the ad hoc models.

Nevertheless, the two sets of variables are strongly related. We have conducted a canonical analysis to further investigate the relationship between the variables used in the economist models and the ones available by the ML algorithms. As expected, since the ML set of inputs includes sub-scores used to construct indicators included in the economist models (e.g. ability), the first canonical correlations are very high (Table A5). A substantial fraction of the higher-order canonical correlations is instead statistically insignificant and less correlated. Therefore, one can conclude that the two sets are closely connected, but that the set of variables used in the ML algorithms might have included additional powerful predictors.

A third, and final set of explanations for why the machine learning did not perform better are related to the machine learning implementation. Section 4.4 already noted that the tuning parameters were chosen by cross-validation and are in the interiors of the ranges considered, so that poor fit is unlikely to be due to tuning. There are a wide variety of different machine learning methods, which raises the possibility of whether using other methods would perform better. Different machine learning algorithms may capture different features of the data. For instance, SVR allows flexible functional forms, while Boosting can include multiple high-order interactions among variables. Furthermore, human experts and algorithms may provide complementary predictions. Therefore, combining these methods may potentially lead to superior performances. Nevertheless, ensemble models merging predictions from LASSO, SVR and Boosting does not lead to substantial improvements (Table A18). Similar results are also obtained when constructing ensembles merging predictions from the machine learning algorithms with the Raven univariate regression, the business plan scores, or the Fafchamps and Woodruff (2017) model. In addition to this, Table A19 shows that using an alternative commonly-used machine learning algorithm, Elastic Net, leads to out-of-sample performances similar to those in Tables 4-6 and A10.

5.3. Predicting high performance

Fig. 1 showed substantial variation in outcomes for both winners and non-winners, with the performance of the very best firms much better than those of the typical firm. For example, the average number of workers among the non-winning firms in the top 10% of employment is 16 (median 13), while the average in the bottom 90% is 2 (median 1, i.e. no employees) and the average profits in the top decile is almost three times that in the rest of the firms. Even if the human experts and machine predictions have difficulty distinguishing between firms across the whole distribution, the question remains as to whether they fare better when trying to identify in advance these very top performers.

To investigate this, we use the predictions obtained in the previous sections and reported in Tables 5 and 6. For each model, we sort firms based on the predicted third-year employment or profit levels. We generated a dummy variable equal to one if the firm is in the top 10%, zero otherwise. Using the actual employment and profit levels, we also generate for each outcome a dummy variable equal to one if the firm is

⁹ In contrast, if we used data from a standalone firm panel survey to predict firm growth, there would be a concern that any relationships found in the data might disappear if the data were collected under competition conditions.

indeed in the top 10%, zero otherwise. The accuracy rate is then obtained by computing the proportion of firms in the hold-out sample correctly predicted to be in the top decile or the bottom 90%.¹⁰ Since we are particularly interested in being able to detect which firms will excel, we also compute the recall rate – i.e. the proportion of the top tail of firms that were correctly predicted.¹¹

Panels A and B of Table 7 reports the results for the top tail in terms of employment and profits respectively. We see that the using the highest business plan scores given by judges only correctly identifies 13 to 16 percent of the firms that are in this top tail. The Raven test, ad hoc models, and machine learning models show slightly better performance, typically identifying around 20 percent of the firms that end up in the top tail, with this reaching as high as one-third for the LASSO in identifying tail employment, and 25 percent for SVM and the Raven score in identifying the top tail of profits. However, note that the sample size for these calculations is relatively small – since we are trying to predict which 21 firms will be best in a hold-out sample of around 210 firms in each case. The result is that accuracy confidence intervals for all the different methods overlap.

From an investor viewpoint, it not only matters whether or not they can identify the exact firms that end up in the top 10%, but how bad their mistakes are if they choose incorrectly. In panel C of Table 7, we therefore consider the case of an investor who is picking the top 10% of firms. We assume this investor makes an investment in each of these firms, and then would get a percentage of profits, so that their goal is to maximize total profits of the firms they choose. For each model, we then aggregate up the profits of the firms chosen in panel B, to report the total monthly profits that the investor would achieve if choosing firms randomly, or according to each selection method. For the non-winners, using the business plans scores from judges would yield returns that are twice as high as random choice, and using any of the Raven test score, Fafchamps and Woodruff model, and machine learning models would yield three to four times the profits. For the sample of competition winners, using the business plan scores would actually yield a portfolio earning lower profits than choosing firms by pure random chance, using the economist models or machine learning models would give approximately twice the level of profits, and the Raven test score yields four times the profits. These results provide suggestive evidence that an investor could use these modeling approaches to gain higher returns, but the bootstrapped confidence intervals are wide and overlap each other.

6. Conclusions and discussion

Our results from a large business plan competition show that successful entrepreneurial performance is very difficult to predict. Expert judges who know the local context, models from several economists with vast experience studying entrepreneurship in developing countries, and multiple machine learning algorithms all struggle to predict which entrants to the competition will succeed most from among those who reached the stage of submitting business plans. We find judges' scores are uncorrelated with firm outcomes three years later. This does not mean business performance is completely unpredictable, as several basic demographic characteristics of the entrepreneurs (age and gender) and

¹⁰ We have followed this approach of setting a fixed number of businesses to predict as successful instead of the alternative of predicting one (success) if the predicted probability is above 0.5 (or the mean of the outcome variable, 0.1 in our case) because we deem this more realistic and similar to how a venture capitalist would act when having to decide on which firms to invest.

¹¹ An advantage of our machine learning approach is the flexibility to consider different objective functions. For instance, since each start-up may have different requests in terms of initial capital necessary to found or expand the business, some investors may actually be interested in maximizing the recall rate subject to a fixed amount of money invested in the businesses predicted to be successful. The calibration procedure can be easily adapted to take this objective into account (see also Sansone, 2018).

measures of their ability (Raven test score) do increase prediction accuracy: males in their 30s who score highly on an ability test are more likely to succeed. However, the overall predictive power of simple models of economists which use these variables is still low, and machine learning algorithms that consider many more potential inputs do not typically outperform these simple models.¹² Even when machine learning algorithms do perform better than the other models considered, the improvements are not substantial and the confidence intervals are large and overlapping.

These results suggest several implications for economists studying entrepreneurship, and for governments and investors seeking to identify high-return entrepreneurs. First, although entrepreneurship is difficult to predict, economists do appear to be capturing the most relevant characteristics for predicting performance, and adding local expert judgement or machine learning does not offer systematic performance gains. Second, for policymakers running competitions, expensive and time-consuming scoring by judges may be able to be replaced by less costly simple predictive models. Third, despite large heterogeneity in outcomes, the fundamental riskiness inherent in entrepreneurship will make it difficult for investors to determine who will grow most, potentially restricting the flow of capital to high-return entrepreneurs.

Several caveats are worth noting. First, our analysis is all conditional on individuals applying for a business plan competition and getting through the first round to the point of submitting a business plan. This process likely selects out many individuals with very low growth prospects, and it may be easier to predict who definitely will not succeed than to predict who might. For example, de Mel et al. (2010) find that measures of background, ability, and attitudes can help distinguish subsistence self-employed from those likely to hire workers. Nevertheless, the problem facing investors and judges of business plan competitions is to decide among those who have self-selected into participating and passed through initial screening, and we show that this is hard to do, despite substantial variation in business performance among the sample.

Second, the poor performance of machine learning in our context may even be an upper bound on how well one could expect such an approach to work in other business plan competitions. We have data from the world's largest business plan competition, and use a training set of outcomes for these participants to predict out of sample fit in a test sample that comes from the same country, population, and time period. We might expect fit to be even worse if models were trained on one competition and then applied to predicting success in another. Policy makers and investors would typically have to make investment decisions based on the past performance of a more limited number of firms. However, our ad hoc economist models of McKenzie (2015) and Fafchamps and Woodruff (2017) did draw upon what the authors had learned from reviewing data from several sources and studying findings from similar contexts. The machine learning algorithms did not have access to these previous datasets, thus limiting by construction their ability to learn, so it is possible that combining data from multiple business plan competitions in different countries could further improve performance.

Third, although we have used three of the most popular machine learning algorithms, it is possible that more advanced algorithms or extensive grid-searchers could further improve performance. However, such algorithms may still be computationally infeasible, even for big companies (Johnston, 2012), or extremely difficult to code, which is the reason behind the extremely high prizes – often reaching \$1 million (Netflix, 2009) – offered in machine learning competitions. Often these

¹² This lack of improvement from machine learning is in line with results in several other domains. For example, Beattie et al. (2018) found a simple average of seven non-academic variables performed as well as sophisticated machine learning techniques when analyzing the importance of personality traits in college student performance; while a simple autoregressive model has similar performances to Google's Flu trends algorithm (Goel et al., 2010).

Table 7
Prediction (out-of-sample) accuracy for tail.

Panel A: Tail Total Employment									
Model	Predictors	Non-winners				Winners			
		Accuracy		Recall	Accuracy		Recall		
		Mean	C.I.		Mean	C.I.			
1	OLS	Constant	79.2%	[74.9%, 83.4%]	7.4%	82.9%	[79.2%, 86.7%]	10.5%	
2	OLS	Business Score	81.0%	[76.5%, 85.5%]	14.8%	83.9%	[79.8%, 87.9%]	15.8%	
3	OLS	Raven Test	81.9%	[77.6%, 86.3%]	18.5%	84.8%	[81.0%, 88.7%]	21.1%	
4	OLS	McKenzie	81.0%	[76.4%, 85.7%]	14.8%	84.8%	[80.9%, 88.8%]	21.1%	
5	OLS	FaWo	81.9%	[77.5%, 86.4%]	18.5%	84.8%	[80.8%, 88.8%]	21.1%	
6	LASSO	All	85.6%	[81.4%, 89.9%]	33.3%	87.7%	[83.7%, 91.7%]	36.8%	
7	SVM	All	81.9%	[77.4%, 86.4%]	18.5%	83.9%	[80.0%, 87.7%]	15.8%	
8	Boosting	All	83.8%	[79.1%, 88.5%]	25.9%	88.6%	[84.8%, 92.5%]	42.1%	

Panel B: Tail Monthly Profits									
Model	Predictors	Non-winners				Winners			
		Accuracy		Recall	Accuracy		Recall		
		Mean	C.I.		Mean	C.I.			
1	OLS	Constant	79.2%	[74.6%, 83.9%]	12.9%	80.3%	[76.2%, 84.3%]	8.3%	
2	OLS	Business Score	80.2%	[75.3%, 85.1%]	16.1%	81.3%	[76.7%, 85.8%]	12.5%	
3	OLS	Raven Test	81.1%	[76.4%, 85.9%]	19.4%	84.1%	[79.8%, 88.5%]	25.0%	
4	OLS	McKenzie	80.2%	[75.5%, 84.9%]	16.1%	81.3%	[76.9%, 85.6%]	12.5%	
5	OLS	FaWo	81.1%	[76.3%, 85.9%]	19.4%	80.3%	[75.9%, 84.7%]	8.3%	
6	LASSO	All	81.1%	[76.2%, 86.0%]	19.4%	80.3%	[76.1%, 84.5%]	8.3%	
7	SVM	All	83.0%	[78.3%, 87.8%]	25.8%	83.2%	[78.8%, 87.5%]	20.8%	
8	Boosting	All	80.2%	[75.1%, 85.3%]	16.1%	81.3%	[77.0%, 85.5%]	12.5%	

Panel C: How much difference would this make for an investor taking a share in the top firms?						
Model	Predictors	Non-winners			Winners	
		Mean	C.I.	Mean	C.I.	
1	OLS	Constant	2.04	[-3.26, 7.33]	4.76	[-7.64, 17.15]
2	OLS	Business Score	4.88	[1.44, 8.32]	4.08	[1.14, 7.02]
3	OLS	Raven Test	5.96	[1.39, 10.54]	19.39	[-2.36, 41.14]
4	OLS	McKenzie	5.40	[0.01, 10.79]	7.92	[-3.03, 18.86]
5	OLS	FaWo	6.74	[0.76, 12.71]	7.26	[-1.78, 16.29]
6	LASSO	All	6.07	[1.72, 10.41]	7.98	[-0.04, 15.99]
7	SVR	All	7.66	[3.05, 12.27]	10.58	[2.33, 18.82]
8	Boosting	All	6.76	[0.79, 12.74]	8.42	[0.20, 16.64]

Notes: outcome in Panel A is whether a firm is in the top 10% for total employment three years after applying (employment coded as zero for applicants not operating firms). Outcome in Panel B is whether a firm is in the top 10% for monthly profits three years after applying (profits coded as zero for applicants not operating firms). Models are estimating separately for competition non-winners and winners. Firms are sorted based on the employment or profit levels predicted in Tables 5 and 6 respectively, and the top 10% are assigned value one, zero otherwise. Out-of-sample accuracy is the ratio of true positives (1) and true negatives (0) to all observations, and is computed using the 20% hold-out sample. The numbers in brackets are bootstrapped 95 percent confidence intervals for hold-out prediction performance. The recall rate is the percentage of businesses correctly predicted to be in the top 10% over the total number of businesses in the 20% hold-out sample actually in the top 10%. Profits in Panel C are in 1,000,000 Nigerian Naira (1 Nigerian Naira equals around 0.0028 US\$). For each model, we have taken the predicted profits computed in Table 6, sort firms in the hold-out sample based on these predicted values, selected the top 21 firms (equivalent to around 10% of the hold-out sample), and summed the profits from these selected firms.

algorithms have been designed to be applied to large data sets with millions of observations, and the sample sizes of firm samples may limit the gains to be had from such an approach.

Finally, our focus has been on predicting which firms will have the best outcomes, which is an object of most interest to investors and lenders. In contrast, such predictions are usually inappropriate for a policy maker aiming to optimally allocate resources (Athey, 2017), who would instead be most interested in allocating funds to firms that would benefit most from the grants, not to those firms which would be successful anyway.¹³ But even in such cases, identifying a pool of firms that are likely to survive and grow may be a useful first step in screening applications, and policy makers may also care about offering support to the next flagship company.

¹³ See Hussam et al. (2016) for a recent application of using machine learning to predict which small entrepreneurs will have the highest returns to grants.

Appendix A. Supplementary data

Supplementary data related to this article can be found at <https://doi.org/10.1016/j.jdeveco.2019.07.002>.

References

- Astebro, T., Elhedhli, S., 2006. The effectiveness of simple decision heuristics: forecasting commercial success for early-stage ventures. *Manag. Sci.* 52, 395–409.
- Athey, S., 2017. Beyond prediction: using big data for policy problems. *Sci. Magna* 355, 483–485.
- Aulck, L., Velagapudi, N., Blumenstock, J., West, J., 2016. Predicting Student Dropout in Higher Education. ArXivWorking Pap arXiv 1606.06364.
- Banerjee, A.V., Duflo, E., 2014. (Dis)Organization and success in an economics MOOC. *Am. Econ. Rev.* 104, 514–518.
- Bauer, E., Kohavi, R., Chan, P., Stolfo, S., Wolpert, D., 1999. An empirical comparison of voting classification algorithms: bagging, boosting, and variants. *Mach. Learn.* 36, 105–139.
- Beattie, G., Laliberte, J.-W.P., Oreopoulos, P., 2018. Thrivers and divers: using non-academic measures to predict college success and failure. *Econ. Educ. Rev.* 62, 170–182.
- Belzon, S., Chatterji, A.K., Daley, B., 2017. Eponymous entrepreneurs. *Am. Econ. Rev.* 107, 1638–1655.

- Blanchflower, D.G., Levine, P.B., Zimmerman, D.J., 2003. Discrimination in the small-business credit market. *Rev. Econ. Stat.* 85, 930–943.
- Boden, R.J., Nucci, A.R., 2000. On the survival prospects of men's and women's new business ventures. *J. Bus. Ventur.* 15, 347–362.
- Bruhn, M., 2009. Female-Owned Firms in Latin America: Characteristics, Performance, and Obstacles to Growth. *World Bank Policy Res. Work. Pap.* p. 5122.
- Celiku, B., Kraay, A., 2017. Predicting Conflict. *World Bank Policy Res. Work. Pap.* p. 8075.
- Chalfin, A., Danieli, O., Hillis, A., Jelveh, Z., Luca, M., Ludwig, J., Mullainathan, S., 2016. Productivity and selection of human capital with machine learning. *Am. Econ. Rev. Pap. Proc.* 106, 124–127.
- de Mel, S., McKenzie, D., Woodruff, C., 2010. Who are the microenterprise owners?: evidence from Sri Lanka on tokman v. de Soto. In: Lerner, J., Schoar, A. (Eds.), *International Differences in Entrepreneurship*. University of Chicago Press, pp. 63–87.
- Duckworth, A.L., Peterson, C., Matthews, M.D., Kelly, D.R., 2007. Grit: perseverance and passion for long-term goals. *J. Personal. Soc. Psychol.* 92, 1087–1101.
- Fafchamps, M., Woodruff, C., 2017. Identifying gazelles: expert panels vs. Surveys as a means to identify firms with rapid growth potential. *World Bank Econ. Rev.* 31, 670–686.
- Friedman, J., 2001. Greedy function approximation: a gradient boosting machine. *Ann. Stat.* 29, 1189–1232.
- Friedman, J., Hastie, T., Tibshirani, R., 2000. Additive logistic regression: a statistical view of boosting. *Ann. Stat.* 28, 337–407.
- Goel, S., Hofman, J.M., Lahaie, S., Pennock, D.M., Watts, D.J., 2010. Predicting consumer behavior with Web search. *Proc. Natl. Acad. Sci.* 107, 17486–17490.
- Guenther, N., Schonlau, M., 2016. Support vector machines. *STATA J.* 16, 917–937.
- Guzman, J., Stern, S., 2017. Nowcasting and placecasting entrepreneurial quality and performance. In: Haltiwanger, J., Hurst, E., Miranda, J., Schoar, A. (Eds.), *Measuring Entrepreneurial Businesses: Current Knowledge and Challenges*. University of Chicago Press, pp. 63–109.
- Hall, R.E., Woodward, S.E., 2010. The burden of the nondiversifiable risk of entrepreneurship. *Am. Econ. Rev.* 100, 1163–1194.
- Hastie, T., Tibshirani, R., Friedman, J., 2009. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer Series in Statistics, second ed. Springer.
- Hussam, R., Rigol, N., Roth, B., 2016. Targeting High Ability Entrepreneurs Using Community Information: Mechanism Design in the Field. *Work (Pap)*.
- Hvide, H.K., Panos, G.A., 2014. Risk tolerance and entrepreneurship. *J. Financ. Econ.* 111, 200–223.
- Jean, N., Burke, M., Xie, M., Davis, W.M., Lobell, D.B., Ermon, S., 2016. Combining satellite imagery and machine learning to predict poverty. *Sci. Magna* 353, 790–794.
- Johnston, C., 2012. Netflix Never Used its \$1 Million Algorithm Due to Engineering Costs (Wired).
- Kerr, W.R., Lerner, J., Schoar, A., 2014. The consequences of entrepreneurial finance: evidence from angel financings. *Rev. Financ. Stud.* 27, 20–55.
- Kleinberg, J., Lakkaraj, H., Leskovec, J., Ludwig, J., Mullainathan, S., 2017. Human decisions and machine predictions. *Q. J. Econ.* 133, 237–293.
- Kleinberg, J., Ludwig, J., Mullainathan, S., Obermeyer, Z., 2015. Prediction policy problems. *Am. Econ. Rev. Pap. Proc.* 105, 491–495.
- Kopf, D., 2015. What Happened when Nigeria Created the World's Largest Business Plan Competition [WWW Document]. *Priconomics*.
- Luca, M., Kleinberg, J., Mullainathan, S., 2016. Algorithms need managers, too. *Harv. Bus. Rev.* 104, 96–101.
- Maroco, J., Silva, D., Rodrigues, A., Guerreiro, M., Santana, I., de Mendonça, A., 2011. Data mining methods in the prediction of Dementia: a real-data comparison of the accuracy, sensitivity and specificity of linear discriminant analysis, logistic regression, neural networks, support vector machines, classification trees and random forests. *BMC Res. Notes* 4, 299.
- McKenzie, D., 2018. Can business owners form accurate counterfactuals? Eliciting treatment and control beliefs about their outcomes in the alternative treatment status. *J. Bus. Econ. Stat.* 36, 714–722.
- McKenzie, D., 2017. Identifying and spurring high-growth entrepreneurship: experimental evidence from a business plan competition. *Am. Econ. Rev.* 107, 2278–2307.
- McKenzie, D., 2015. Identifying and Spurring High-Growth Entrepreneurship: Experimental Evidence from a Business Plan Competition. *World Bank Policy Res. Work. Pap.* p. 7391.
- McKenzie, D., Paffhausen, A.L., 2018. Small firm death in developing countries. *Rev. Econ. Stat.*
- Michelacci, C., Schivardi, F., 2016. Are They All like Bill, Mark, and Steve? the Education Premium for Entrepreneurs. *EIEF Work.* p. 16. *Pap.*
- Mullainathan, S., Spiess, J., 2017. Machine learning: an applied econometric approach. *J. Econ. Perspect.* 31, 87–106.
- Nanda, R., 2016. Financing high-potential entrepreneurship. *IZA World Labor.* 1–10.
- Netflix, 2009. *Netflix Prize* [WWW Document]. URL: <http://www.netflixprize.com/> (accessed 1.1.17).
- Ng, A., 2016. *Machine Learning* [WWW Document]. Coursera. URL: <https://www.coursera.org/learn/machine-learning> (accessed 1.1.16).
- Nikolova, E., Ricka, F., Simroth, D., 2011. Entrepreneurship in the transition region: an analysis based on the Life in Transition Survey. In: *Crisis and Transition: the People's Perspective*. EBRD, pp. 76–96.
- OECD, 2010. *High-Growth Enterprises: what Governments Can Do to Make A Difference*. OECD Stud. SMEs Entrep.
- Olafsen, E., Cook, P.A., 2016. *Growth Entrepreneurship in Developing Countries: A Preliminary Literature Review*. The World Bank Group, Washington, D.C.
- Queiro, F., 2016. The effect of manager education on firm growth. *Q. J. Econ.* 118, 1169–1208.
- Reddy, S., 2017. A Computer Was Asked to Predict Which Start-Ups Would Be Successful. The Results Were Astonishing. *World Econ. Forum*.
- Ricadela, A., 2009. *Fifty Best Tech Startups*. Bloom. *Businessweek*.
- Robb, A.M., Watson, J., 2012. Gender differences in firm performance: evidence from new ventures in the United States. *J. Bus. Ventur.* 27, 544–558.
- Sansone, D., 2018. Beyond early warning indicators: high school dropout and machine learning. *Oxf. Bull. Econ. Stat.*
- Schonlau, M., 2005. Boosted regression (boosting): an introductory tutorial and a Stata plugin. *STATA J.* 5, 330–354.
- Scott, E., Shu, P., Lubynsky, R., 2015. Are “better” ideas more likely to succeed? An empirical analysis of startup evaluation. *Harvard Bus. Sch. Work. Pap.* 16.
- Skriabikova, O.J., Dohmen, T., Kriechele, B., 2014. New evidence on the relationship between risk attitudes and self-employment. *Labour Econ.* 30, 176–184.
- Subrahmanian, V.S., Kumar, S., 2017. Predicting human behavior: the next frontiers. *Sci. Magna* 355, 489.
- van Praag, C.M., Cramer, J.S., 2001. The roots of entrepreneurship and labour demand: individual ability and low risk aversion. *Economica* 68, 45–62.
- Varian, H.R., 2014. Big data: new tricks for econometrics. *J. Econ. Perspect.* 28, 3–28.
- Verplancke, T., Van Looy, S., Benoit, D., Vansteelandt, S., Depuydt, P., De Turck, F., Decruyenaere, J., 2008. Support vector machine versus logistic regression modeling for prediction of hospital mortality in critically ill patients with haematological malignancies. *BMC Med. Inf. Decis. Mak.* 8, 56.
- Zacharakis, A.L., Shepherd, D.A., 2001. The nature of information and overconfidence on venture capitalists' decision making. *J. Bus. Ventur.* 16, 311–332.
- Zygmunt, Z., 2016. What is better: gradient-boosted trees, or a random forest? [WWW Document]. *FastML*. URL (accessed 1.1.17). <http://fastml.com/what-is-better-gradient-boosted-trees-or-random-forest/>.