

It's a Peoples Game, Isn't It?! A Comparison Between the Investment Returns of Business Angels and Machine Learning Algorithms

Entrepreneurship Theory and

Practice

00(0) 1–38

© The Author(s) 2020



Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/1042258720945206

journals.sagepub.com/home/etp

Ivo Blohm^{1*} , Torben Antretter^{1*}, Charlotta Sirén^{1,2} ,
Dietmar Grichnik¹ , and Joakim Wincent^{1,3}

Abstract

Investors increasingly use machine learning (ML) algorithms to support their early stage investment decisions. However, it remains unclear if algorithms can make better investment decisions and if so, why. Building on behavioral decision theory, our study compares the investment returns of an algorithm with those of 255 business angels (BAs) investing via an angel investment platform. We explore the influence of human biases and experience on BAs' returns and find that investors only outperformed the algorithm when they had extensive investment experience and managed to suppress their cognitive biases. These results offer novel insights into the role of cognitive limitations, experience, and the use of algorithms in early stage investing.

Keywords

business angels, artificial intelligence, machine learning, biases, investment experience, decision making

Business angels (BAs) play a vital role in financing the growth of early stage ventures (Carpentier & Suret, 2015; Maxwell et al., 2011). Between 2014 and 2016, BAs invested €18.5 billion in more than 100,000 deals, accounting for 67% of early stage investments in Europe (EBAN, 2017). Despite their intense investment activity, BAs can only process a limited proportion of the input they receive and must therefore base many of their investment decisions on incomplete and uncertain information (Harrison et al., 2015; Maxwell et al., 2011; Payne et al., 1993). Although prior research has emphasized these cognitive limitations, including decision biases, scholars

¹University of St.Gallen, St.Gallen, Switzerland

²University of Queensland, UQ Business School, Brisbane Qld, Australia

³Hanken School of Economics, Helsinki, Finland

*Drs Ivo Blohm and Torben Antretter contributed equally to the work.

Corresponding Author:

Charlotta Sirén, Global Center for Entrepreneurship & Innovation, University of St.Gallen, St.Gallen, Switzerland.

Email: charlotta.siren@unisg.ch

have not yet considered their impact on BAs' investment performance nor alternatives that could go beyond the limitations of the human mind in making early stage investment decisions.

The rise of artificial intelligence (AI) and machine learning (ML) algorithms has enabled not only fast but also highly accurate information processing beyond the scope of human abilities (Hastie et al., 2009; Jordan & Mitchell, 2015) thus providing fruitful opportunities for decision making in entrepreneurship (e.g., Obschonka & Audretsch, 2019; Schwab & Zhang, 2019). In particular, recent advancements in ML point to the fact that unlike human decision makers, algorithms can consistently process large amounts of unstructured and rapidly changing data from an almost indefinite variety of sources to make unbiased predictions about future events (Agarwal & Dhar, 2014; Blattberg & Hoch, 1990). Therefore, these algorithms are increasingly being adopted in the financial markets (Wigglesworth, 2016). Even though research has not yet shown the usefulness of ML in BA investing, such technologies will likely become increasingly relevant now that the market for early stage venture finance is developing towards more competitive and complex funding trajectories (Bessière et al., 2020). The current trend is that BAs are forming larger and more visible membership groups and investment platforms—both offline and online—that increase their visibility and exposure to deal flow but often require them to make decisions even faster (Bessière et al., 2020; Bonini et al., 2018; Drover et al., 2017). In this highly dynamic and complex environment, ML algorithms present a promising path to a potential new era in early stage investing.

The objective of this study is to compare the investment returns of BAs who are prone to decision biases with those of ML algorithms on the same investment opportunities. Behavioral decision theory suggests that decision makers' experience plays a critical role in decision quality and speed, especially under conditions of extreme uncertainty (Tversky & Kahneman, 1974). Similarly, in the BA context, studies have shown that with increasing investment experience (i.e., number of investments made), BAs develop mental shortcuts that allow them to reach decisions much faster (Harrison et al., 2015; Maxwell et al., 2011). For instance, Wiltbank et al. (2009) provided evidence that investor experience has a positive effect on investment returns. Above that, Huang (2018) showed that expert intuition accurately predicts which investments will be extraordinarily profitable. However, a recent study by Capizzi (2015) showed that the relationship between BAs' investment experience and investment returns may be more complex and that it is important to consider whether experts are less sensitive to cognitive biases in decision making. Thus, even though heuristics and intuition, honed by experience, might have benefits in BA investing, this research area warrants further attention.

This article addresses three important research questions: (1) Do the investment returns of ML algorithms outperform those of BAs? (2) If they do, is this due to different human decision biases? (3) Are highly experienced BAs less sensitive to decision biases and thus able to generate higher investment returns than ML algorithms when investing in early stage ventures? To answer these questions, we compare the investment returns of a state-of-the-art ML algorithm with the investment returns of 255 BAs who invested via a large angel investment platform. The angel investment platform in our study provides BAs with a standardized set of information about each investment opportunity. This information includes key facts (e.g., industry, year of founding, social media profiles), a brief self-descriptive text, and a presentation document about the focal venture. Thus, this platform is an ideal setting for conducting such a comparison because it allows us to use the exact same information that was available to the BAs to train our ML algorithm. Above that, we can use data sources like LinkedIn to quantify soft information about entrepreneurs (e.g., founder human capital) and compare the trained algorithm's assessments to those of BAs who "rely on their gut feel, rather than formal analyses" (Huang & Pearce, 2015) to evaluate the quality of a founder's team. While the decision accuracy of humans is likely to decline with increasing amounts of unstructured data, the decision accuracy

of ML algorithms is likely to enhance with additional data inputs. Compared to the general notion that the machine always wins (Witten et al., 2016), our context of BA investing is characterized by extreme uncertainty that embraces expert intuition (Huang, 2018) and thus questions the potential superiority of AI when compared to experienced BAs that manage to suppress their decision biases.

We draw on theory on behavioral decision making and the role of human experience (e.g., Tversky & Kahneman, 1974) to test the assumption regarding BAs' cognitive limitations and the extent to which experience is helpful in BA investing (Harrison et al., 2015). By comparing the outcomes of BAs' early stage investment decisions to those of a ML algorithm, our study makes four contributions. First, it adds to the literature on BA decision making (e.g., Carpentier & Suret, 2015; Mason & Harrison, 1996; Maxwell et al., 2011) and the emerging body of research on their investment returns (e.g., Capizzi, 2015; Mason & Harrison, 2002; Wiltbank et al., 2009; Wiltbank, 2005) by showing that, on average, human investors have lower investment returns than their artificial counterparts. This result suggests that the value of machine-based decision support systems to improve early stage investment returns.

Second, we show that BAs' decision biases are the reason algorithms generally outperform them. We outline how common decision biases—i.e., local bias, overconfidence, and loss aversion—impact the investment returns of BAs who invest via an angel investment platform. Decision biases are often assumed, but it is rarely tested how they can be avoided. Our comparative approach allows us to test the role of individual biases against a mostly unbiased benchmark, thus contributing to the literature on the role of biases and heuristics in early stage investing (e.g., Harrison et al., 2015; Maxwell et al., 2011).

Third, although research has suggested that human decision makers are limited by their information processing abilities due to bounded rationality, we acknowledge that experienced and largely unbiased BAs can still be more successful than ML algorithms. Previous research on experts has contrasted them to novices (Chi et al., 1981; Dew et al., 2009) but has not used the benchmark of an algorithm that is able to process and learn from an indefinite and constantly updated stream of data. Our findings largely support the idea that some experienced BAs who are able to use their experience to suppress their decision biases can even outperform ML algorithms in terms of their investment performance. This implies a selective human advantage in early stage decision making. Therefore, our study adds to our understanding of the boundary conditions in which ML can help improve decision making in highly uncertain environments such as early stage investing. Above that, we show that experienced humans who manage to suppress their biases may still be the “gold standard” in using affective judgment to recognize patterns in uncertain decision environments. For investors who do not belong to this selected group of BAs, ML may provide substantial benefits in making investment decisions.

Fourth, researchers have just recently started to embrace the possibility of using ML in entrepreneurship research (e.g., Obschonka & Audretsch, 2019; Schwab & Zhang, 2019). Our study follows the emerging trend of leveraging ML approaches to investigate important but previously unstudied areas in entrepreneurship (see Obschonka & Audretsch, 2019 for a summary of relevant articles) and is—to the best of our knowledge—the first study to compare the investment returns of an algorithm to those of actual investors in early stage venture finance. Our study outlines how entrepreneurship scholars can use a “transfer learning” approach, i.e., transferring knowledge from a domain with abundant training data to domains in which it is difficult to train accurate ML algorithms due to the lack of data, to study entrepreneurial phenomena. We thus contribute to the dialogue on the usefulness of ML for early stage investing and entrepreneurship research.

Theoretical Background and Hypotheses

Business Angels' Information Processing and Machine Learning in Early Stage Investment Decisions

Investment decisions in entrepreneurial ventures are often described as complex due to the large quantities of unstructured and uncertain information (Huang, 2018). Thus, individuals' ability to process information and make reliable decisions is key to explaining the outcomes of angel investments (Harrison et al., 2015; Maxwell et al., 2011). According to behavioral decision theory, processing information inputs reliably is difficult for two main reasons: human investors' bounded rationality (Simon, 1955) and their cognitive biases (Tversky & Kahneman, 1974).

First, due to bounded rationality, BAs have limits in their cognitive capacity to process information, which decreases their ability to make rational investment decisions that take every potential piece of information into account (Maxwell et al., 2011; Tversky & Kahneman, 1974). Therefore, BAs must capture and manage multiple signals and cues to make sense of and mentally account for the complex considerations of an early stage investment decision (Huang, 2018). These arguments suggest that BAs might suffer from "analysis paralysis" (Huang, 2018, p. 1824) that influences their ability to conduct thorough assessments and make the best possible investment decision, which, in turn, negatively impacts their investment returns.

Second, behavioral decision theory argues that cognitive biases cause decision makers to process information incorrectly, which may lead to inaccurate judgments that impact the quality of their decisions (Tversky & Kahneman, 1974), especially when it comes to interpreting soft information (Huang, 2018; Huang & Pearce, 2015). In the entrepreneurial finance literature, many biases, such as local bias, overconfidence, and loss aversion have been suggested to influence BAs' decision making (Maxwell et al., 2011).

Local bias is a well-documented bias that describes an investors' preference for local investment targets (Coval & Moskowitz, 1999; Harrison et al., 2010). Local investments are defined as those that reside within the same region or nation as the investor (Coval & Moskowitz, 1999; Harrison et al., 2010; Jääskeläinen & Maula, 2014). The literature on proximity preference and local bias suggests that the root cause of local bias is that investors struggle to identify and evaluate investment opportunities that are distant from their own environment (Jääskeläinen & Maula, 2014). The lack of information regarding distant investment opportunities may lead BAs to either ignore or mistrust the completeness and trustworthiness of information on investment opportunities (Hirshleifer, 2001; Huberman, 2001). In the venture capital (VC) context, several studies have shown how local bias affects the composition of investment portfolios (e.g., Cumming & Dai, 2010; Jääskeläinen & Maula, 2014; Sorenson & Stuart, 2001). Recently, Hornuf and Schmitt (2016) showed that local bias also exists in equity crowdfunding. Interestingly, their findings indicate that BA-like investors have even stronger local bias than other types of investors in equity crowdfunding. Furthermore, Sorenson and Stuart (2001) showed that local bias also exists among investors who intend to play a rather passive role in their portfolio companies. Based on these arguments and empirical evidence, we expect local bias to influence BAs' decision making when evaluating early stage investment opportunities.

Overconfidence is one of the most prevalent decision making biases. Accordingly, research has asserted that "no problem in judgment and decision making is more prevalent and more potentially catastrophic than overconfidence" (Plous, 1993, p. 217). Overconfidence is the tendency to overestimate one's knowledge and/or investment skills and, consequently, the likelihood of a set of events (e.g., Li & Tang, 2010; Moore & Healy, 2008; Navis & Ozbek, 2016). For example, in the context of early stage investments, overconfident investors may overestimate the likelihood that a company will succeed (Zacharakis & Shepherd, 2001). Although overconfidence itself does not necessarily

lead to wrong decisions, this bias is likely to prevent learning processes (Zacharakis & Shepherd, 2001). In other words, overconfident BAs may not fully consider all relevant information nor search for additional information to improve their investment decisions. In the VC context, Zacharakis and Shepherd (2001) empirically showed that the majority of VC investors are overconfident and that overconfidence negatively influences their decision accuracy. Interestingly, they showed that the more information is available to investors the worse their decisions are because additional information increases their overconfidence even further. We believe the same logic applies to the BA investment context, especially to angel investment platforms, where BAs have plenty of information on numerous investment opportunities.

Loss aversion refers to the fact that humans tend to be more sensitive to potential losses (e.g., loss of wealth) than to potential gains (Thaler et al., 1997). Empirical studies have found that losses are weighted about twice as much as gains (e.g., Kahneman et al., 1990; Tversky & Kahneman, 1992). This insight is especially relevant for BAs who can only invest in a finite number of companies, when suffering from a loss can significantly impact their overall portfolio returns (Mason & Harrison, 2002). As such, Benjamin and Margulis (1996, p. 221) observed that it is often “much more important [for BAs] to avoid a bad investment than to try to hit a home run.” Another example of loss aversion is what Cumming et al. (2005) call liquidity risk. BAs may fear not being able to effectively exit an investment and thus being forced to remain invested much longer or sell their investment at a high discount. Taken together, we expect loss aversion to influence BAs returns from their investment decisions in an unfavorable manner.

In contrast to human decision makers’ biased decisions, making predictions about future events using ML has become more accessible in recent years (George et al., 2016). Given their capacity to process large amounts of data, ML algorithms may lead to results that are unbiased, quickly produced, and free of social or affective contingency. Moreover, ML algorithms do not suffer from cognitive resource limitations when processing and interpreting information from a variety of sources (Blattberg & Hoch, 1990). As opposed to BAs, who employ “elimination-by-aspect” decision models (Maxwell et al., 2011), ML algorithms can weigh the value of all available information when processing large amounts of complex data and can thus make more integral investment decisions that may lead to superior investment performance. Furthermore, ML algorithms are mostly described as unbiased (Hernandez et al., 2019). As such, they are able to pick companies globally (i.e., avoid local bias), which improves their set of available investment opportunities compared to BAs. Above that, the non-existence of overconfidence in ML algorithms enables them to consider all available information and thus avoid misjudgments in their decision making processes. Finally, algorithms may have an advantage in investing as they can be set up to weigh gains and losses equally and thus freely choose among the most promising investment opportunities.

Although—to the best of our knowledge—no study has compared the results of ML algorithms with the investment returns of early stage investors, research from the VC field has provided promising insights into algorithms’ ability to predict the outcomes of later stage investments. For instance, Krishna et al. (2016) used financial metrics from Crunchbase, such as the number of investment rounds, to accurately predict the success and failure of 11,000 ventures. Investigating VC-backed ventures, Hunter et al. (2017) constructed hypothetical investments and, using random draws, showed that the exit rates of their models’ portfolios were as high as 60% and thus much higher than the conventional exit rates of VC portfolios, as reported by, for instance, Hochberg et al. (2007). These empirical insights further support our theoretical notions that machine intelligence may also outcompete human investment decisions in early stage ventures, thus providing a means to generate higher investment returns. Given these arguments, we hypothesize the following:

Hypothesis 1a: Early stage investments selected by ML algorithms are more likely to generate high- investment returns than those selected by business angels.

Hypothesis 1b: The difference in the investment performance between ML algorithms and business angels stems from human decision biases.

The Role of Business Angels' Investment Experience in the Comparison Between Business Angels and Machine Learning Algorithms

The literature on experts' decision making has shown that, with experience, decision makers develop intuition-based decision models to aid them in making decisions for which the risks are highly uncertain (Newell et al., 1958). Using intuition-based decision models distilled from repeated investment experience has proven to be a valuable strategy for making early stage investment decisions (Huang & Pearce, 2015). The intuition-based decision making model proposes that experienced BAs may use their affective judgments, or intuition, to overcome the limitations of their bounded rationality such as their "analytics overwhelm" (Huang & Pearce, 2015, p. 658), and thus be able to effectively predict which investments will be successful (Huang, 2018; Huang & Pearce, 2015). Experts are generally described as being talented in making intuitive judgments on information that statistical models might not be able to capture (Einhorn, 1974). Since algorithms depend on "means-end calculation and reliance on abstract and universally valid rules" (Lindebaum et al., 2020, p. 16), they would most likely discard some opportunities that expert investors identify as highly attractive. Experience has also been shown to alleviate the influence of various decision biases in investment-related decision making (e.g., Dimov et al., 2007). We expect that expert BAs (i.e., those who have gained substantial investment experience), as opposed to their fellow investors with less experience, are able to make use of an affective judgment tool—namely, their intuition—to disentangle the complexity of early stage investment decisions. Thus, we expect that experts are less sensitive to the influence of decision biases and will, as a consequence, perform equally well or even better than unbiased algorithms.

When comparing experts' and novices' decisions under conditions of uncertainty, research has posited that experts adopt a different decision making logic than novices do and that their decisions thus strongly differ from each other (Kahneman & Tversky, 1973; Tversky & Kahneman, 1974). Novices and experts rely on different heuristics, or decision making strategies, to account for complex or incomplete information when evaluating alternatives and acting upon opportunities. There is evidence that investors' experience is important (e.g., Capizzi, 2015; Collewaert & Manigart, 2016; Wiltbank et al., 2009) because accumulated experience increases the efficiency of mental shortcuts (i.e., expert heuristics; Holcomb et al., 2009). Expert heuristics, or knowledge structures, are created by tacit knowledge that allows individuals to impose a structure to make decisions using emerging and ambiguous information. As such, related research has found that contextual wisdom results in faster and more accurate problem solving in specific knowledge domains (Dew et al., 2009; Mitchell et al., 2009; Simon & Simon, 1978). Further, experience enables individuals to apply analytical reasoning to go beyond what is evident in their short-term memory (the actual information provided). Thus, experience allows individuals to craft theories from small quantities of data, integrate potential contingencies, and make decisions that surpass the patterns provided by the information they have at hand, all without the "analytics overwhelm" (Huang & Pearce, 2015). Above that, experts can possess highly specialized domain knowledge that provides them with a competitive advantage in terms of recognizing and interpreting rare information.

The literature on debiasing suggests that experience leads to learned rational behavior and allows decision makers to mitigate or even avoid decision biases (e.g., Chen et al., 2007; Sorenson & Stuart, 2001). The underlying argument in early stage investing is that BAs previous experience allows them to focus on the most important areas in analyzing an investment opportunity while ignoring others, thus suppressing their decision biases (Collewaert & Manigart, 2016; Dimov et al., 2007). For instance, in the context of local bias, Cumming and Dai (2010) found that investors use their experience and network to overcome information asymmetry stemming from the geographical distance to their investments (see also Hsu, 2004). Furthermore, Sorenson and Stuart (2001) showed that experience alters the influence of spatial distance (i.e., local bias) on investment decisions because investors become less sensitive to distance as they accrue investment experience. Increased investment experience thus expands BAs information networks and increases the geographical reach of their investments (Sorenson & Stuart, 2001). The evidence regarding the relationship between overconfidence and experience is mixed. While some scholars have shown that overconfidence might increase with experience (e.g., Heath & Tversky, 1991; Shane, 2009), others have argued that investors' overconfidence, although it might increase with initial experience, should decrease with higher investment experience because investors may become better at recognizing their own abilities (e.g., Gervais & Odean, 2001). In the BA context, this contemporary view on overconfidence would suggest that with more investment experience, BAs develop better self-assessments (Gervais & Odean, 2001), which help them make more deliberate investment decisions. Huang and Pearce (2015) showed that expert BAs' intuition-based heuristics accurately predict which investments would be extraordinarily profitable. In other words, expert BAs may make investment decisions that are counterintuitive but provide significant earnings potential (i.e., choosing positive outliers), which reflects the opposite logic of loss aversion. Overall, experience may provide investors with a valuable toolset to mitigate or even overcome the influence of decision biases. Therefore, we expect that experienced BAs are able to generate higher investment returns than their inexperienced counterparts when investing in early stage ventures.

Because experienced BAs are able to mitigate the influence of cognitive biases, they are able to achieve similar levels of objectivity in their decision making as ML algorithms that are unemotional and mostly free of decision biases. Although ML algorithms have the same information on which to base a decision, including soft information (e.g., founding team composition, communication), we expect that experienced (largely debiased) BAs will generate better investment decisions than rational machines under conditions of extreme uncertainty. The main reason lies in the human investor's emotional and social intelligence to understand, contextualize, and read the weak signals of soft information (i.e., the "human touch" [Raisch & Krakowski, in press, p. 16]). Research on expert BA decision making often describes their "gut feel" about the entrepreneur and the success chances of the venture as a key criterion and driver of their investment decisions (Hisrich & Jankowicz, 1990; Huang, 2018; Huang & Pearce, 2015). Although ML algorithms can be exposed to soft types of information, they do not possess human senses, perceptions, emotions, and social skills (Braga & Logan, 2017), which are needed, for example, to persuade others to support an investment decision. Above that, ML algorithms are not able to understand fundamental technical and ethical challenges because they are limited in interpreting physiological reactions and social norms that humans typically follow (McDuff & Czerwinski, 2018; Raisch & Krakowski, in press). Based on these arguments, we expect that experienced and largely debiased BAs will generate better investment decisions than rational ML algorithms. Taken together, we hypothesize that:

Hypothesis 2a: *When comparing experienced with inexperienced BAs, experienced BAs are more likely to generate higher investment returns than inexperienced BAs when investing in early stage ventures.*

Hypothesis 2b: *When comparing experienced BAs with ML algorithms, experienced BAs with lower levels of decision biases are likely to generate higher investment returns than ML algorithms when investing in early stage ventures.*

Methods

Study Context and Dataset

We test our hypotheses in the context of an angel investment platform. Angel investment platforms enable BAs to make faster investment decisions with less administrative and financial effort than BAs acting individually (e.g., Bonini et al., 2018; Drover et al., 2017). Above that, angel investment platforms expose BAs to increased deal flow with standardized information about investment opportunities (Bonini et al., 2018; Croce et al., 2017). This information often includes self-descriptive texts about the venture's business model, management team, and the market as well as references to other sources of information, such as the venture's website or social media accounts.

The angel investment platform examined in our study is one of the largest in Europe and has experienced steady membership growth over the past years.¹ During our study period—the period in which we were able to extract data from the platforms deal monitoring system (i.e., December 2013 to June 2018)—the number of BAs investing via the platform rose from 136 in 2013 to 472 in June 2018. Of the 472 BAs, 363 made at least one investment, and 255 made three or more investments. Because our primary goal was to analyze investment portfolio returns, we followed Mitteness et al. (2016) and defined that a portfolio must consist of at least three portfolio companies.

The majority of the 255 BAs with three or more investments in their portfolios reside in Switzerland (17%), Belgium (15%), France (9%), the United States (7%), and Germany (6%). The platform itself does not make investments. Instead, each BA decides for him or herself whether to participate in an investment opportunity. When a deal is submitted via the platform's website, it is pre-screened by trained staff members from the platform to ensure it meets the members' overall investment criteria (e.g., stage of development or investment size). Once a deal has passed the pre-screening it is added to the platform for the investors' consideration. If a group of BAs is interested in making an investment, the deal is assigned to an experienced deal leader (i.e., a fellow BA from the network), who functions as the contact person for the venture and leads the negotiation on behalf of her or his fellow investors. Once the deal is made, the deal leader usually takes a board seat with the company, while the other BAs who co-invested in the deal (average 17.36 BAs per investment; standard deviation (SD): 21.10) remain passive. The BA group members hold regular in-person meetings to discuss the current status, view new pitches, and make decisions about specific investment opportunities.

Our sample of investment opportunities contains data on 623 pre-screened early stage ventures that sought funding on the angel investment platform and were invited to provide a live pitch to the network members.² The companies from our set of investment opportunities were founded between 2005 and 2017 (average: May 2013; SD: 2.34 years). By the end of our study period in June 2018, our sample BAs invested in 51 individual ventures from 14 different countries, with the majority of investments being made in Switzerland (58%), France (12%), and Belgium (8%). The main industries in which the BAs invested are information and communications technology (32%) and industrial products (20%). The average deal size on the platform was

€127,704 (SD: €192,459; average in Europe: €184,000; EBAN, 2017), with each of the 255 BAs investing on average €9,528 per investment (SD: €9,867). Based on the ventures' net asset values (NAV), which are calculated by the platforms administration twice a year in line with *International Private Equity and Venture Capital Valuation Guidelines* (IPEV, 2018), 76% of the BAs' investment portfolios in our sample had a positive internal rate of return (IRR).

The strength of our dataset lies in the detailed information on the investor and deal levels. On the investor level, we were able to obtain direct cash flow information on each investment from the platform's deal monitoring system. This feature was crucial for our prediction approach because it enabled us to calculate precise measures of investment performance and benchmark the returns of our hypothetical ML portfolios against real investors. Above that, our sample included standardized information on the ventures, including their names, names of founders, industry tags, social media accounts, etc. We used the information on the venture level to mine additional data from different publicly available data sources such as Twitter, LinkedIn, Google, and official national trade registers. We used available application programming interfaces (API) to generate large-scale information—to the extent permitted by each APIs terms of use—from these data sources. Data sources that did not allow for programmatic data mining were mined manually. As a result, we mined 422,444 tweets, 11,913 LinkedIn entries and 757 Google Trends Index keywords.

Prediction Approach

As the volume and variety of available data have rapidly increased in the internet era, analyzing (big) data exceeds the capability of many known research methods in entrepreneurship (Obschonka & Audretsch, 2019; Schwab & Zhang, 2019; van Witteloostuijn & Kolkman, 2019). As such, there is a growing need for algorithms to structure, analyze, and model data. The term ML broadly refers to algorithms that optimize performance criteria by evaluating predicted outputs against observed (or “true”) data points.³ Typically, ML algorithms have the capacity to uncover non-obvious patterns in data and facilitate reliable and accurate predictions about future events. Compared to classical regression approaches that focus on explaining the variance in a given dataset and aim for causality, ML algorithms intend to predict future outcomes. As such, ML algorithms may accept a lower explanatory power in exchange for the ability to perform more accurate predictions that can be generalized to new data.

In our study, we used gradient boosted decision trees, a state-of-the-art ML method (Chen & Guestrin, 2016; Friedman, 2001; Hastie et al., 2009), to make our predictions. Gradient boosted decision trees have recently been applied to predict new venture survival (Antretter et al., 2018, Antretter et al., 2019), bank failure (Carmona et al., 2019; Climent et al., 2019), credit risk assessments (Chang et al., 2018), and cryptocurrency fluctuations (Li et al., 2019). The approach is based on the idea of creating multiple decision trees, with each tree segmenting data into groups by following a series of decision rules (James et al., 2013). In order to make a prediction for a given observation, each tree makes a decision on which group the observation belongs to and then uses the group's average to make a prediction. The general idea of boosting is to combine the predictions of many “weak” decision trees to produce a powerful “ensemble” that provides more accurate predictions (Hastie et al., 2009). In this approach, each new decision tree is trained to improve the existing ensemble of decision trees by learning from their past prediction errors (Chen & Guestrin, 2016). In other words, newly trained decision trees are fitted to the residuals from the combined ensemble of trees and not the outcome itself. Thus, new decision trees, which can consist of very simple individual models, are integrated into the ensemble so that they update the residuals until a certain minimum is reached (James et al., 2013). In our study, we used “eXtreme Gradient Boosted Trees” (XGBoost) by Chen and Guestrin (2016),

which is currently one of the most advanced implementations of gradient boosted decision trees. For a more detailed elaboration on our XGBoost algorithm, see Appendix 1.

In order to compare the investment performance of our XGBoost predictions to BAs' returns, we followed a three step process: (1) data mining, (2) algorithm training, and (3) comparing investment performance.

(1) Data Mining

To increase the breadth and depth of our baseline dataset (i.e., 623 ventures that were invited to pitch their company), we mined publicly available data from various online data sources. First, we followed Raz and Gloor (2007) and performed an automated company search on the internet to obtain information on whether the company was still operating. Second, we collected data from Twitter, LinkedIn, Google Trends, and national trade registers to enrich each venture's data profile.

We focused our data mining on publicly available data sources that BAs could obtain for free in order to inform their decision making processes. To prevent issues of target leakage,⁴ we used data that could be related to the date of founding (e.g., founders' experiences and education before founding the focal venture). In some selected cases, that is, for Twitter data, the relevant data has been created by default after this point of time (e.g., tweets posted by a venture). In this case, we mined the data created in the first year of operation only. We then use these digital traces to make predictions beyond this initial year. As such, we were able to create robust predictions because we timewise disentangle each predictor from the dependent variable of the prediction. All data was mined in June 2018.

(2) Algorithm Training

Our dataset of 623 investment opportunities only provided performance data for the 51 ventures the BAs actually invested in. Therefore, we trained the XGBoost algorithm to predict new venture survival, which we were able to determine for all 623 ventures. We made this decision for both theoretical and methodological reasons. First, from a theoretical point of view, survival is often described as the most relevant and reliable performance measure for early stage ventures (see Soto-Simeone et al., 2020 for a review). Determining NAVs or IRRs without records like term sheets, capitalization tables, or audited financial statements could be considered highly subjective. As such, survival has been the go-to measure for many studies investigating early stage investment performance (e.g., Dimov & De Clercq, 2006; Manigart et al., 2002). Second, from a methodological point of view, training a robust ML algorithm requires a certain number of cases. Since our research partner—one of the largest BA investment platforms in Europe—has only invested in 51 early stage companies at the time of our study, building a robust algorithm with these cases alone would not have been possible. Thus, using survival as an alternative allowed us to utilize the entire dataset for our analysis. This strategy is inspired by the concept of “transfer learning,” which is regularly used in the ML domain (Pan & Yang, 2009). The concept refers to the attempt of transferring knowledge from a domain with abundant training data (e.g., new venture survival) to a domain in which it is difficult to train accurate ML models due to the lack of training data (e.g., early stage ventures' financial performance).

As a first step to train the algorithm, we partitioned the dataset of 623 ventures into two separate parts (James et al., 2013). We built the model with one part (i.e., training data) and evaluated its predictive accuracy by applying it to an unused part of the data (i.e., testing data). This separation is important because gradient boosted decision trees tend to “memorize” all patterns found in the training data due to the sequential updating process. As a consequence, their prediction accuracy might be very high on the training data. On the unseen testing data, however, their prediction performance may drop significantly (a phenomenon called “overfitting”). To get a

realistic understanding of our algorithm's ability to predict unseen data that has not been part of training the algorithm, we used the 51 actual investments as the testing set and the remaining 572 ventures as the training set.

Second, to mitigate potential sampling bias resulting from using data of one specific angel investment platform, we applied stratified bootstrapping (Bickel & Freedman, 1984). Bootstrapping relies on the logic of resampling from a dataset to approximate the actual distribution of parameters (Efron & Tibshirani, 1994). In total, we performed 2,500 bootstrap iterations. In each bootstrap iteration, we randomly drew subsamples with replacement (i.e., 400 observations; ca. 70% of training data) from our training set to build XGBoost models that predict new venture survival on different parts of the data. Following the original distribution of companies in the training data that had not survived, we stratified all bootstrapped samples so that they matched the original distribution. The results of our approach are very robust to different sizes of the bootstrapped samples, as well as the number of bootstrap iterations.

(3) Comparing Investment Performance

To compare the (hypothetical) investment performance of our algorithm with the actual returns of the BA's, we followed the ideas of Hunter et al. (2017). First, we predicted survival, or, more specifically, the risk of dying, for each of the 51 ventures using our XGBoost algorithm. We evaluated the predictive performance of each model with Harrell's concordance index.⁵ Second, we built hypothetical investment portfolios by using the predicted risks of dying as the only decision criteria. As such, we made investment choices based on the idea that the algorithm would always pick the ventures with the lowest estimated risk of dying. For instance, for a portfolio of three ventures, the algorithm would choose the three ventures with the lowest risk of dying from a given set of investment opportunities. Based on these hypothetical investment choices, we determined the corresponding IRRs for each investment choice and averaged them on the portfolio level. Using the trained XGBoost models, we constructed portfolios of different sizes. Following our sampling for the BA portfolios, we created portfolios that contained hypothetical investments between 3 and 15 portfolio companies. This is important because we matched randomly sampled BA portfolios of equal sizes to each of these hypothetical portfolios as counterfactuals and compared their average IRRs with those of the algorithm (Hunter et al., 2017). In doing so, we were able to prevent our results from being biased by influential outliers that would be overly used in our sampling procedure. We verified that our results are stable for different portfolio cutoffs.

Measures

In the following, we describe the measures (i.e., predictors) that we used to train our XGBoost algorithm. Table 1 provides a detailed overview of all the predictors used in our study including their operationalizations, descriptive statistics, and relative importance for the predictive model.

Survival

We determined new venture survival by conducting an automated company search on the internet to investigate whether each venture had a website and, if so, whether the website suggested that the company was still alive (for a similar approach see Raz & Gloor, 2007). When the company's website returned an error, was blank, or displayed a notice indicating that the new venture went out of business, we inferred that the company was no longer active. Using the internet archiving tool Wayback Machine,⁷ which regularly stores versions of publicly accessible websites, we inferred each of the focal ventures' time of death. The Wayback Machine has been used extensively in management and entrepreneurship research to investigate archived website

Table 1. Variables and Descriptive Statistics.⁶

Variable	Operationalization	Data Source	Reference	Mean	SD	Min	Max	Importance Rank
Online legitimacy								
Negative emotions (tweets)	Avg. negative emotional valence of the tweets posted by a venture in the first year.	Twitter API	Liew and Wang (2016)	1.31	1.48	0.00	14.22	1
No. of tweets	Total no. of tweets posted by a new venture in the first year.	Twitter API	Shepherd and Zacharakis (2003)	120.18	238.32	0.00	2,161	5
Length of tweets	Avg. length of characters of a new venture's tweets of the first year.	Twitter API	Shepherd and Zacharakis (2003)	87.85	43.45	0.00	235.00	6
Positive emotions (retweet)	Avg. positive emotional valence of first year's tweets for all languages.	Twitter API	Liew and Wang (2016)	3.16	4.27	0.00	41.77	7
Positive emotions (tweets)	Avg. positive emotional valence of first year's tweets for all languages.	Twitter API	Liew and Wang (2016)	3.87	3.65	0.00	39.48	10
No. of retweets	No. of retweets performed by a new venture in the first year.	Twitter API	Shepherd and Zacharakis (2003)	8,723	146,383	0.00	3,389,606	11
Negative emotions (replies)	Avg. negative emotional valence of the replies given by a venture in the first year.	Twitter API	Liew and Wang (2016)	0.74	1.66	0.00	17.03	15
No. of replies	No. of replies given by the new venture in the first year.	Twitter API	Liew and Wang (2016)	12.77	37.03	0.00	295.00	16
Positive emotions (replies)	Avg. positive emotional valence of the replies given by a venture in the first year.	Twitter API	Liew and Wang (2016)	4.38	7.62	0.00	50.00	20
No. of tweet languages	Total no. of languages used in first year's tweets.	Twitter API	Shepherd and Zacharakis (2003)	3.00	2.84	0.00	19.00	25
Negative emotions (retweet)	Avg. negative emotional valence of first year's tweets for all languages.	Twitter API	Liew and Wang (2016)	1.41	3.26	0.00	39.83	26

(Continued)

Table 1. Continued

Variable	Operationalization	Data Source	Reference	Mean	SD	Min	Max	Importance Rank
Initial Funding and Legal Form								
Equity capital before BA funding	Obtained initial funding at the time of the founding in Euros.	National trade registers		74,757	152,383	0.00	2,140,000	2
Legal form	Dummy (=1) if legal form is other, 0 = corporation.	National trade registers		0.42	0.49	0.00	1.00	18
Market Timing								
Google Trends Index industry	Standardized Google Trends Index the industry a venture is operating in at the time of the founding.	Google Trends Index		45.98	24.88	-3.39	92.82	4
Google Trends Index product	Standardized Google Trends Index for the tags describing the venture's product/service at the time of the founding.	Google Trends Index		41.33	13.69	8.32	90.39	8
Year of founding	Year in which the new venture was founded.	BA platform		2012	2.35	2005	2017	24
Market Timing								
B2C/B2B	Dummy (=1) if B2C business model, 0 otherwise (=B2C).	BA platform	Böhm et al. (2017)	0.52	0.50	0.00	1.00	3
Founder Experiences								
Managerial experience	Sum of months founders worked in a management job before founding.	LinkedIn	Baptista et al. (2014)	92.22	123.72	0.00	754.33	9
Technical experience	Sum of months founders worked in a technical job before founding.	LinkedIn	Hunter et al. (2017)	35.63	63.63	0.00	582.00	14
Entrepreneurial experience	Sum of months founders worked in their own company before founding.	LinkedIn	Westhead and Wright (1998)	21.87	47.57	0.00	346.67	21
Founder Experiences								
Consulting experience	Sum of months founders worked in top-tier consulting before founding.	LinkedIn	Antretter et al. (2018)	4.57	20.59	0.00	250.00	31

(Continued)

Table 1. Continued

Variable	Operationalization	Data Source	Reference	Mean	SD	Min	Max	Importance Rank
Investment banking experience	Sum of months the founders worked in investment banking before founding.	LinkedIn	Shepherd and Zacharakis (2003)	4.49	24.39	0.00	350.50	36
Industry								
AgriTech/FoodTech	Dummy if industry is AgriTech/ FoodTech (=1) or not (=0).	BA platform		0.01	0.12	0.00	1.00	12
ICT	Dummy if industry is Information and Communication Technology (=1) or not (=0).	BA platform		0.36	0.48	0.00	1.00	13
Consumer Goods	Dummy if industry is Consumer Goods (=1) or not (=0).	BA platform		0.11	0.31	0.00	1.00	17
Mobility/TravelTech	Dummy if industry is Mobility/ TravelTech (=1) or not (=0).	BA platform		0.03	0.17	0.00	1.00	19
FinTech/InsureTech	Dummy if industry is FinTech/ InsureTech (=1) or not (=0).	BA platform		0.12	0.33	0.00	1.00	23
BioTech	Dummy if industry is BioTech (=1) or not (=0).	BA platform		0.02	0.12	0.00	1.00	27
CleanTech/GreenTech	Dummy if industry is CleanTech/ GreenTech (=1) or not (=0).	BA platform		0.03	0.18	0.00	1.00	29
HighTech	Dummy if industry is HighTech (=1) or not (=0).	BA platform		0.07	0.25	0.00	1.00	30
Other Industry	Dummy if industry is "other industry" (=1) or not (=0).	BA platform		0.12	0.32	0.00	1.00	33
MedTech/HealthTech	Dummy if industry is MedTech/ HealthTech (=1) or not (=0).	BA platform		0.06	0.23	0.00	1.00	34
RegTech/LegalTech	Dummy if industry is RegTech/ LegalTech (=1) or not (=0).	BA platform		0.01	0.10	0.00	1.00	40
PropTech	Dummy if industry is PropTech (=1) or not (=0).	BA platform		0.01	0.12	0.00	1.00	42
Founder Education								
MBA degree	Sum of founding members who have a MBA degree.	LinkedIn	Gimeno et al. (1997)	0.19	0.44	0.00	2.00	22

(Continued)

Table 1. Continued

Variable	Operationalization	Data Source	Reference	Mean	SD	Min	Max	Importance Rank
Bachelor's degree	Sum of founding members who have a bachelor's degree.	LinkedIn	Gimeno et al. (1997)	0.59	0.78	0.00	6.00	28
Master's degree	Sum of founding members who have a master's degree.	LinkedIn	Cooper et al. (1991)	0.81	0.99	0.00	8.00	32
Uni top tech	Avg. no. of founding members who attended a top technical university.	LinkedIn, Times Higher Education	Antretter et al. (2018)	0.15	0.44	0.00	3.00	35
Business education	Avg. number of founders that completed business education.	LinkedIn	Antretter et al. (2018)	0.79	0.84	0.00	6.00	37
Technical education	Avg. number of founders that received a technical education.	LinkedIn	Antretter et al. (2018)	0.76	0.91	0.00	7.00	38
PhD degree	Sum of founding members who have a PhD degree.	LinkedIn	Cooper et al. (1991)	0.21	0.51	0.00	3.00	39
Uni top business	Avg. no. of founding members who attended a top business university.	LinkedIn, Times Higher Education	Antretter et al. (2018)	0.25	0.53	0.00	3.00	41

content and infer website age (e.g., Funk, 2014; Scott et al., 2020; Walthoff-Borm et al., 2018; Wry et al., 2014). If a venture was still alive, we inferred that a new venture was surviving until the end of our study period (i.e., right-censoring data), which is a common approach in survival analysis. The data was measured on a monthly basis. For each venture, we collected the time span between the time of founding and the time of death/end of the study period in months (see Appendix 1 for details).

Survival Predictors

In order to predict the new venture's risk of dying, we included a series of predictors that have been shown to influence new venture survival: (1) legitimacy, (2) human capital, (3) market timing, and (4) funding. First, legitimacy has been argued to be one of the most important factors in understanding new venture survival (e.g., Shepherd & Zacharakis, 2003; Zimmerman & Zeitz, 2002). As such, a recent prediction study has found that online legitimacy can be used to predict new ventures' survival (e.g., Antretter et al., 2019). Second, entrepreneurs' human capital has a major influence on the survival prospects of new ventures (Delmar & Shane, 2006 for a review). For instance, Gimeno et al. (1997) showed that having supervised others (i.e., having management experience) before starting a new venture has a positive impact on survival. Third, there is a growing strand of research arguing that the market timing of products and business models has a strong impact on new venture survival (e.g., Boyer & Blazy, 2014; Hyytinen et al., 2015). Finally, initial capital is described as a vital resource associated with new venture survival. A higher amount of capital permits strategies that are more ambitious and provides a more significant margin of error (Cooper et al., 1991). As such, many scholars have found initial capital to be positively associated with new ventures' chances of survival (e.g., Cooper et al., 1991; Headd, 2003; van Praag, 2003).

(1) Legitimacy. To operationalize legitimacy, we followed previous research (e.g., Antretter et al., 2019; Etter et al., 2018) and analyzed the ventures' Twitter accounts. The "Twitter Account Activity API" allowed us to access all tweets that were posted in the first year of operation (Kuppuswamy & Bayus, 2017). In order to measure the ventures' online legitimacy, we focused on quantity of information and information content (Antretter et al., 2019).

Quantity of information. The quantity of information about a venture's products, management team, and organization may be considered a main source of legitimacy because it provides a frame of reference for potential customers to understand the benefits of the venture's offerings (Shepherd & Zacharakis, 2003). As such, we measured the number of tweets that were posted in the first year, the number of languages that were used in these tweets, and the tweets' average length.

Information content. To evaluate the content of the ventures' communication, we measured the emotional valence of tweets, retweets, and replies. To do so, we used the Linguistic Inquiry and Word Count (LIWC), developed by Pennebaker et al. (2001), which has been used extensively in entrepreneurship research (e.g., Fischer & Reuber, 2011; Parhankangas & Renko, 2017; Wolfe & Shepherd, 2015). LIWC is a dictionary-based text mining approach in which a text is compared against a dictionary (e.g., containing words associated with positive emotions) in order to determine how strongly a text is associated with the constructs that are operationalized by the dictionary. In our study, we used the positive and negative emotion LIWC dictionaries in English, German, French, Italian, Spanish, Dutch, and Portuguese (Pennebaker et al., 2001) to determine the sum of positive or negative words in each tweet that were normalized by the tweet's length. These "term frequencies" can thus be considered as a standardized measure of positive and negative emotions for each tweet. Also, we analyzed emojis in the ventures' tweets by translating them into their English description⁸ and applying the same LIWC analysis.

All variables for measuring online legitimacy were created at the level of the individual tweet and were aggregated to the venture level by using the arithmetic mean.

(2) Human Capital. To approximate the founders' human capital, we analyzed the LinkedIn profiles for all founder teams of the ventures. Unfortunately, LIWC does not provide experience- and education-related dictionaries and, to the best of our knowledge, no other validated sources for measuring human capital from textual data exists. Thus, we analyzed the founders' LinkedIn profiles in a LIWC-like fashion using a set of novel dictionaries to determine the founders' prior work experiences and formal education. Below, we shortly describe the obtained measures. Appendix 2 reports technical details, shows the dictionaries, and describes their development and validation process.

Prior work experience. Based on the founders' LinkedIn profiles, we analyzed the founders' past positions in terms of whether or not they conveyed managerial, technical, or entrepreneurial experience. Also, we investigated whether the founders worked in a top management consultancy or investment bank because these types of jobs have shown to yield a specific set of skills (e.g., problem-solving or communication) that may help entrepreneurs succeed with their businesses (Dimov & Shepherd, 2005; Zarutskic, 2010). Applying our dictionaries to the founders' LinkedIn profiles, we could determine whether a given position in a founder's LinkedIn profile was associated with one of these types of professional experiences. Triangulating this information with the reported duration of the position, we were able to quantify the obtained experience in months. If a given position contained more than one type of experience (e.g., a position describing a Chief Technology Officer might contain managerial and technical experience), the duration of the position was equally split among the involved experience types. All variables were measured in terms of total duration (i.e., months) and then aggregated by the arithmetic mean to the venture level. For our prediction approach, we only used data until the date the focal venture was founded.

Education. We inferred the founders' level of formal education (Gimeno et al., 1997) from their LinkedIn profiles. More specifically, we measured whether the founders had a PhD, MBA, master's degree, bachelor's degree, or did not possess a college degree. These variables reflected dummies for which "no college degree" served as the reference group. Above that, we differentiated whether founders had received a technical or managerial education and whether founders had been enrolled in either a leading technical or business university (according to the "Times Higher Education Ranking").

(3) Business Model, Industry, and Market Timing. We determined each venture's business model, the industry it operates in, and a proxy for its market timing to include standardized information about each venture's product or service into our prediction approach.

Business model. Based on the ventures' description on the angel investment platform, we differentiated between business-to-customer (B2C) and business-to-business (B2B) models by including a dummy variable that takes the value of 1 for B2C and 0 for B2B.

Industry. We used the BA platform's industry classification to determine which industry each venture belonged to. Industries were dummy-coded with "Ad & RetailTech" serving as a randomly selected reference group. A list of the investigated industries is provided in Table 1.

Market timing. All ventures described themselves with a series of keywords ("tags") on the angel investment platform (e.g., "artificial intelligence" or "virtual reality"). We used these keywords to proxy the market timing of the ventures and their products and/or services. We followed previous research (e.g., Yuan & Lee, 2018) and used the Google Trends Index—a standardized coefficient of the amount of Google search requests for a given keyword over time—to measure market timing. As such, we constructed two measures for each venture based on the Google Trend Index. First, we determined the Google Trends Index for each industry (e.g., "FinTech") at the time the venture was founded. Second, we used the self-descriptive keywords provided by the venture (e.g., "artificial intelligence") to

determine their Google Trends Index at the time of the founding. Multiple keywords for an investment opportunity were averaged by arithmetic mean. The rationale behind these measures was to determine whether a venture was founded in times of a market hype (e.g., indicated by a high Google Trend Index at the time of the founding) or whether it could be considered an early mover (e.g., with a low Google Trend Index). Measures for the different keywords were standardized. Further, we tracked the year of founding for each venture as indicated on the angel investing platform.

(4) Equity Capital Before BA Funding. To measure the initial funding the ventures received when they started, we investigated all national trade registers and related sources from the ventures' country of origin to determine their equity capital in Euros at the time of the founding (i.e., before the company received its first funding on the platform). To control for the effect of different legal requirements, we also mined data related to the ventures' legal entity. More specifically, we measured whether the venture was founded as a corporation or another legal form. The data was coded as a dummy with "corporation" serving as the reference group.

Internal Rate of Return

We calculated IRR to compare the investment performance of the BAs and the hypothetical investments of our prediction approach. For the BA portfolios, we followed Mason and Harrison (2002) and calculated IRRs based on the NAVs of each BA investment taking into account both positive and negative returns (Capizzi, 2015) as well as the holding period over which valuation changes occurred (see Equation 1). We derived the share prices at different points in time (i.e., on the purchase date and at the end of our study period) from the platform's deal monitoring system. We determined the holding period in years between these time points to calculate annualized IRRs. We aggregated these investment-level IRRs to the portfolio level by volume-weighted average. Following Mason and Harrison (2002), we did not account for any income from dividends or fees, e.g., for taking board positions, that the BAs may have received. For ventures where the BAs invested in multiple rounds, we aggregated the IRRs across all rounds by volume-weighted mean. We matched these IRRs to the investment choices of our algorithm to mimic the performance of the hypothetical portfolios.

$$0 = \sum_{n=0}^N \frac{CF_n}{(1 + IRR)^n} = CF_0 + \frac{CF_N}{(1 + IRR)^N} \quad (1)$$

where:

CF_0 = Share price at the initial investment

CF_n = Share price at period n

CF_N = Final share price

N = Holding period, that is, years between initial investment
and final evaluation of investment

n = Period

IRR = Internal rate of return

Business Angel's Investment Experience

We followed Capizzi (2015) and measured investment experience by computing the number of investments (i.e., individual portfolio companies) the BAs made via the platform (see also Collewaert & Manigart, 2016).

Business Angel's Decision Biases

To assess the role of BAs' cognitive limitations, we measured (1) local bias, (2) overconfidence, and (3) loss aversion.

(1) Local bias. Local bias describes BAs' tendency to make investments in close proximity to their own location (Harrison et al., 2010; Mason & Harrison, 1994; Wetzel, 1981). In line with prior research, a "local" investment is defined as one where the investee resides within the same region or nation as the investor (Coval & Moskowitz, 1999; Harrison et al., 2010; Jääskeläinen & Maula, 2014). We thus measured local bias as the share of investments made in local ventures compared to the total number of portfolio investments each BA made (for a similar approach see van Nieuwerburgh & Veldkamp, 2009).

(2) Overconfidence. Overconfidence describes BAs' tendency to overestimate their knowledge and/or investment skills (Li & Tang, 2010; Moore & Healy, 2008; Navis & Ozbek, 2016). Overconfident investors are more likely to overcommit their limited financial resources to opportunities in which they overestimate their abilities to assess the risk (Hayward et al., 2006; Pikulina et al., 2017; Zacharakis & Shepherd, 2001). As such, we calculated standardized z-scores for the size of each investment. We considered the maximum z-score among all investments in a BA portfolio as a proxy for overconfident investment behavior. Thus, we measured the extent to which an investor makes extreme investments with disproportionately high investment tickets that diverge from her or his regular investment behavior (Pikulina et al., 2017).

(3) Loss aversion. Loss aversion describes BAs' tendency to be more sensitive to potential losses than to potential gains (Thaler et al., 1997). A widely used proxy for the risk in portfolios of venture investors is the fraction of investments made in ventures in the early stages of development (Buchner et al., 2017; Ruhnka & Young, 1991). In his seminal work on the informal venture capital market, Wetzel (1981) showed the expected risk of loss consistently decreases with maturing stages of development. The angel investment platform in our study has categorized each deal as being at seed stage (19% of investments made), early stage (37% of investments made), expansion stage (22% of investments made), or late stage (22% of investments made). For each BA, we thus determined the average investment stage of all portfolio investments by calculating the arithmetic mean across the stages of development with seed stage being coded as 1, early stage as 2, expansion stage as 3, and late stage as 4.

Results

Hypotheses Tests

Table 1 presents the descriptive statistics for all the predictors, including their mean, standard deviation, minimum and maximum, and relative importance (i.e., ranking) for making the predictions.

We used Harrell's Concordance Index to determine the algorithms predictive accuracy. Our XGBoost algorithm reached a Harrell's Concordance Index of 0.60. This index means that 60% of the survival predictions are correct when comparing them across all possible pairs in which at least one venture did not survive (Harrell et al., 1996). The investigated BAs achieved, on average, an IRR of 2.56% (SD: 21.11%). This IRR is higher than the IRRs reported in other studies. For instance, Capizzi (2015) reports an average IRR of 1.78% for a sample of Italian BAs. Our ML algorithm, however, outperformed these returns. On average, our algorithm achieved an IRR of 7.26% (SD: 18.25%). This represents a statistically significant ($p < 0.001$) average

Table 2. Performance of High- and Low-Biased Bas

Type of Bias	High-Biased BAs IRR in %				Low-Biased BAs IRR in %				Significance of Difference
	Mean	SD	Min	Max	Mean	SD	Min	Max	
Local bias	0.00%	16.23%	-45.15%	37.29%	5.14%	24.75%	-55.01%	349.42%	-22.37***
Loss aversion	-0.01%	14.28%	-55.01%	36.57%	5.86%	25.51%	-38.03%	349.42%	-29.90***
Overconfidence	0.01%	16.95%	-55.01%	141.70%	4.22%	24.19%	-36.28%	349.42%	-15.04***
All biases	0.00%	11.94%	-45.15%	9.41%	10.82%	28.85%	-27.86%	349.42%	-31.56***

Notes. $N = 32,500$; *** $p < .001$.

performance gain of 184%. Hypothesis 1a, which proposes that early stage investments selected by ML algorithms are more likely to generate higher investment returns than those selected by BAs, thus receives strong support.

Our data suggest that all three decision biases exist among the BAs. In terms of local bias, 28% of the BAs made at least 50% of their investments within their country of residence; 19% even made more than 75% of their investments within their country of residence. Looking at our overconfidence measure, we see that 91% of the BAs tend to exhibit overconfident investment behavior with selected investments tickets being more than 1 SD higher than their average investment size; 16% even made extreme investments that diverge more than 2 SDs from their average ticket size. Further, our data shows that, on average, only 19% of the investments were made in seed stage ventures, suggesting that most BAs tend to be somewhat loss averse.

Table 2 provides an overview of the returns for BAs that showed high (above the sample median) and low (below the sample median) values for each cognitive bias. It further shows the returns for BAs who had above or below median values for all three decision biases. The results for all individual biases and the combination of biases show that BAs who had lower bias values had higher investment returns compared to BAs with high bias values. All performance differences are highly significant ($p < 0.001$). Thus, our results provide strong support for Hypothesis 1b, confirming that decision biases are an important differentiator of BAs' investment returns and may thus be the reason why unbiased ML algorithms outperform BAs.

Our Hypothesis 2a stated that experienced BAs are more likely to generate higher investment returns than their inexperienced counterparts when investing in early stage ventures. We tested Hypothesis 2a by comparing the IRRs of experienced BAs (above the sample median) to the IRRs of inexperienced BAs (below the sample median). Our results show that the BAs' investment performance generally increased with investment experience. While inexperienced investors achieved an average IRR of 0.31%, more experienced investors achieved an average IRR of 5.20%. This difference is highly significant ($p < 0.001$). Thus, our results provide strong support for Hypothesis 2a.

To further understand the role of investment experience in generating higher investment returns than our ML algorithm and to test Hypothesis 2b, we compared the investment performance of experienced BAs at different levels of the decision biases. For all three biases there is a statistically significant difference between the investment returns of experienced investors with above and below median bias levels (local bias: low bias = 6.80%, high bias = 3.74%, $p < 0.001$; loss aversion: low bias = 14.41%, high bias = -1.34%, $p < 0.001$; overconfidence: low bias = 13.77%, high bias = 2.06%, $p < 0.001$; see Table 3). For further testing Hypothesis 2b, we compared the three groups of experience and unbiased BAs against our ML algorithm. We find that all experienced BAs with low loss aversion and low overconfidence significantly outperform our ML algorithm in terms of IRR ($p < 0.001$). However, we found no statistically significant

Table 3. Performance of Experienced and Inexperienced BAs Across Bias Levels

Level of Experience	Type of Bias	Business Angels IRR in %												Significance of Difference Low Bias BAs vs. ML Algorithm	
		ML Algorithm IRR in %				High Bias				Low Bias					Significance of Difference High vs. Low Bias BAs
		Mean	SD	Min	Max	Mean	SD	Min	Max	Mean	SD	Min	Max		
High	Local bias	7.26%	18.25%	-50.00%	162.73%	3.74%	14.79%	-27.46%	33.13%	6.80%	22.09%	-18.97%	67.06%	-9.84***	-1.67
	Loss aversion					-1.34%	10.25%	-18.97%	9.41%	14.41%	23.41%	-27.46%	67.06%	-49.79***	-22.78***
	Overconfidence					2.06%	13.62%	-27.46%	33.13%	13.77%	26.38%	-18.97%	67.06%	-26.85***	-15.19***
	All biases					2.87%	8.09%	-9.60%	9.41%	22.75%	22.63%	11.19%	67.06%	-47.58***	-37.25***
Low	Local bias					-3.65%	16.72%	-45.15%	37.29%	3.86%	26.55%	-55.01%	349.42%	-22.62***	-11.57***
	Loss aversion					-0.45%	18.10%	-55.01%	36.57%	0.81%	25.36%	-38.03%	349.42%	-3.81***	-24.17***
	Overconfidence					-2.41%	22.79%	-55.01%	141.70%	1.26%	22.66%	-36.28%	349.42%	-9.33***	-26.91***
	All biases					-20.52%	13.04%	-45.15%	-10.95%	3.51%	29.80%	-27.86	349.42%	-39.13***	-8.78***

Notes: N (High experience) = 15,000; N (Low Experience) = 17,500; ***p < .001.




Figure 1. Effect of Local Bias and Experience on Investment Performance

difference between experienced BAs with low local bias and the ML algorithm ($p = n.s.$), that is, these BAs perform on par with the ML algorithm. Thus, Hypothesis 2b is partially supported by our results.

Figures 1–3 illustrate the joint effect of investment experience and decision biases. The figures show that only BAs with high experience and comparatively low bias levels can optimize their investment returns against the unbiased benchmark of a ML algorithm.

Additional Analysis

In order to probe the robustness of our results, we performed a series of additional analyses. First, we investigated our results using different operationalizations of experience. For instance, we used network tenure in months (Mittens et al., 2016). Second, in our main analysis, we compared volume-weighted average IRRs for the BAs with hypothetical portfolios that were aggregated by arithmetic mean. We tested different aggregation approaches for both types of portfolios and have found these approaches to be the best performing aggregation approaches for their respective portfolio types. We also verified that our results and conclusions are robust to different aggregation methods. Third, we calculated risk-adjusted portfolio returns using Sortino ratios, which penalize returns for downside risk (Estada, 2006). Overall, the results are very similar to the results using IRR such that we consider them to be robust to variations of downside risk. Fourth, we investigated whether the set of BAs was influenced by some kind of survival and/or success bias, that is, whether the BAs that made low performance investments discontinued investing. We found no indication that past investment performance is a predictor for future success or investment activity. Fifth, although venture valuations are performed twice a year by the platform's administration, we performed an additional analysis in which we removed all




Figure 2. Effect of Loss Aversion and Experience on Investment Performance

ventures whose valuation did not change since the time of investment. Again, the results led us to similar conclusions. Finally, we investigated whether our approach is also valuable in




Figure 3. Effect of Overconfidence and Experience on Investment Performance

identifying the most profitable ventures among all investment opportunities in our data set (instead of building artificial portfolios of higher performance from a “relevant set” of pre-selected ventures). The results of this robustness test indicate that our ML approach would systematically invest more frequently in ventures with a positive IRR and less frequently in startups with a negative IRR than a random investment process would suggest.⁹

To increase the practical relevance of our results, we further analyzed the importance of each variable for the predictive algorithm. The variable importance ranking shown in Table 1 depicts the contribution of each measure to the overall prediction model. Thus, the higher a variable’s ranking, the higher a variable’s importance for making accurate predictions. The most important variables refer to measures indicating the online legitimacy a new venture possesses (negative sentiment of tweets: ranking first; number of tweets: ranking fifth). The equity capital before BA funding ranks second. Beyond that, the two Google Trends Index variables were important, that is, the market timing of industry tags (ranking fourth) and product tags (ranking eighth). These results suggest the importance of overall market timing in making early stage investment decisions. Furthermore, the business model (B2B vs B2C; ranking third), as well as prior managerial experiences of the founders (ranking ninth) was estimated as being important for making accurate predictions.

Discussion

Theoretical Implications

AI in general and ML algorithms in particular increasingly guide the decisions we make (e.g., Hernandez et al., 2019). Given the abundance of available data from online data sources, many investors have begun to use algorithms to scout promising investment opportunities and help make investment decisions. However, entrepreneurship research has not yet tested the effectiveness of machine-generated investment decisions in an early stage investment context by comparing the investment returns of ML algorithms to those of BAs. Addressing this gap, our study investigates the role of human decision bias and the potential of ML to overcome such biases. Although the literature on AI and ML is rapidly developing, to the best of our knowledge, this is the first study providing evidence of a “man vs. machine” comparison in early stage investing.

By empirically showing that ML algorithms are able to produce, on average, higher investment returns than BAs investing via an angel investment platform, our study contributes to the literature on BA decision making (e.g., Carpentier & Suret, 2015; Mason & Harrison, 1996; Maxwell et al., 2011) and angel investment returns (e.g., Capizzi, 2015; Mason & Harrison, 2002; Wiltbank, 2005; Wiltbank et al., 2009). More specifically, we show that, on average, our ML algorithm is able to achieve a performance gain of up to 184% when compared to the BAs in our sample. The average IRR of 7.26% shown for our ML algorithm is also well above the angel investment returns reported by other studies (e.g., Capizzi, 2015). Furthermore, we disentangle the mechanisms that lead to these findings by investigating the interaction between decision biases, investment experience, and performance. BAs are generally considered to have limited cognitive capacities and fall prey to a series of decision biases, such as local bias, overconfidence, and loss aversion. By showing that the BAs with high levels of decision biases have lower investment returns than those who show lower levels of such biases, we add to the ongoing discussion on biases and heuristics in early stage investing (e.g., Harrison et al., 2015; Huang & Pearce, 2015; Huang, 2018; Maxwell et al., 2011). Algorithms, on the other hand, are not sensitive to these biases and can thus be seen as an optimal benchmark to investigate the role of decision biases in early stage investing (Hernandez et al., 2019).

Building on recent empirical findings and theoretical arguments about the role of learning from experience for achieving higher returns in early stage investing (e.g., Capizzi, 2015; Harrison et al., 2015), we suggest that the superiority of ML algorithms in early stage investing may not be indefinite.¹⁰ We argue that the outcomes of human information processing in generating attractive investment returns under conditions of extreme uncertainty significantly depend on BAs' investment experience. While previous research about angel investing has emphasized the contrast between experts and novices (e.g., Jacoby et al., 2001; Wiltbank et al., 2009), no study has ever compared BAs to ML algorithms that can process an unlimited number of data sources. Our comparison of machine-generated investment portfolios versus BA portfolios shows that although decision biases may lead to inferior investment returns, the quality of decision outcomes increases with more investment experience. Investment experience alone, however, is not enough to produce better investment returns than ML algorithms. Only investors who manage to use their experience to suppress cognitive biases in their decision making are able to produce better returns than our algorithm. This finding highlights an important message on the role of experience in early stage investing (e.g., Capizzi, 2015; Huang & Pearce, 2015; Huang, 2018; Wiltbank et al., 2009): it is a people's game—but only for experienced players. These findings add to the theoretical discussion about cognitive limitations by showing that decision biases, such as local bias, overconfidence, and loss aversion can be mitigated.

Researchers have just begun to investigate the possibilities of ML in entrepreneurship research. We add to the rapidly evolving body of research (see Obschonka & Audretsch, 2019 for a summary of relevant articles) by contributing to the study of the boundary conditions in which ML provides value for early stage investors. Prior research was primed by the notion that an AI (i.e., machines that can act, reason, decide, and solve problems without input limitations) may always be the preferred option (Pomerol, 1997; Schwab & Zhang, 2019). Our study aims to enrich this perspective by providing practically relevant evidence for the applicability of ML algorithms in entrepreneurship. We argue that ML algorithms that are largely unbiased and capable of processing large amounts of unstructured data play an important, complementary role in early stage investment practice. This notion is somewhat contrary to the current discussion in the entrepreneurial finance literature that explicitly emphasizes the value of mental shortcuts and investor gut feeling in early stage investment decision making (e.g., Huang, 2018; Huang & Pearce, 2015; Maxwell et al., 2011), sparking the idea that human gut feel may be irreplaceable. However, these studies have only highlighted the negative aspects of using "large quantities of data [and] multiple types of data" (Huang, 2018) in early stage decision making because humans lack the capacity to process this information effectively. However, with the help of modern technology to mine and structure publicly available data sources that include valuable information about human behavior, emotions, and competencies (i.e., soft information), ML algorithms increasingly provide a means to overcome these limitations of the human mind. The availability of such technologies thus provides fruitful opportunities to improve the overall quality of investment decisions.

Practical Implications

Besides several contributions to the entrepreneurship literature, our study also offers important practical implications for BAs. By analyzing the importance of various factors in making accurate predictions of new venture survival (Table 1), we inform BAs about determinants they should consider when evaluating potential investment opportunities. Our results support the accepted notion that the equity capital before BA funding is an important survival predictor. What is perhaps more surprising is that Twitter-based online legitimacy variables and the Google Trend Index—an indicator for overall market timing—are very important to make accurate

predictions. This implies that social media such as Twitter or other digital traces such as Google Trends could be valuable tools for BAs to get information that may inform their investment decision making. Interestingly, human capital variables such as management or entrepreneurial experience are not as important to make survival-based predictions than many other, more structurally oriented variables such as the business model. These findings also show that well-trained prediction models may be largely constructed from publicly available information, meaning that building these tools is in the reach of individual BAs or at least professionally organized angel investment platforms. Moreover, given that ML algorithms were able to outperform an average BA in our data, our findings recommend BAs consider using ML to make better investment decisions. Although experienced BAs might outperform ML algorithms, collaborating with algorithms to train their decision making and to work more like experts could provide a fruitful opportunity for BAs, as ML algorithms could provide “something to consider as you’re weighing your options” (Hernandez et al., 2019, p. 76). Stock market investing has relied on the support of ML for many years and we believe this is a potential future for BA investing as well. Finally, our study shows early stage investors that asking the question “can this be a viable business that has very high chances of surviving?” might be more valuable in achieving higher portfolio returns than searching for the “needle in the haystack.”

Limitations and Future Research

Our research has some limitations, which suggest avenues for future examination. First, our results suggest that both humans and machines have their distinct advantages and should be integrated to achieve optimal investment performance. We believe that such integration (i.e., approaches of hybrid intelligence) could provide fruitful avenues for future research on performance in BA investing. For example, it might be that the decision logics of ML and BAs complement each other: ML algorithms may operate with decision logics more closely linked to causation, that is a logic of prediction, while expert BAs might resort to effectuation, that is, a logic of control (Saravathy, 2001). In the current study, we are unable to test this idea, as our dataset does not allow us to empirically examine these decision making logics.

Second, although the amount of data we used in our analysis enabled us to provide interesting insights into the potential of ML in early stage investing, we believe that there are great possibilities to generate additional data from online data sources (psychometric profiles, web traffic, etc.), which warrants much more research attention. In addition, generating BA performance data from larger investment pools such as Angel.co or Crunchbase would allow researchers to study BA investments over a longer period of time, thus better accounting for holding periods and exit preferences of BAs.

Third, we use survival data to train our ML algorithm because in the very early stages of development, survival is often the most relevant and accessible performance measure. Although our results were very robust against a series of additional analyses, future research—especially studies on the later stages of venture development—should consider training their algorithms with data that predicts performance directly (e.g., returns, mergers and acquisitions, initial public offering). Above that, in our study, IRRs are determined based on the reevaluation of NAVs. Although our study context (i.e., central evaluations by the platform’s administration instead of BAs evaluating NAVs themselves) mitigates potential issues of self-report bias to a large extent, future studies could aim to investigate performance measures that are solely based on realized returns (e.g., through trade sales or initial public offering).

Fourth, as we generated data from one angel investment platform, our experience measure is limited to the number of investments BAs have made on this specific platform. Although we corroborated our results with alternative experience measures, future research should collect data

that allows for studying BAs' investment experience across their entire career as an investor. In addition, we encourage further studies to explore and distinguish the different learning processes that are linked to experience in an early stage investment setting.

Fifth, although our study included a representative set of decision biases that are well established in the entrepreneurial finance literature, there are plenty of additional biases that are likely to affect investment decisions. Perhaps alternative biases such as herd mentality bias (i.e., BAs' tendency to follow and copy other BAs) or confirmation bias (i.e., BAs' tendency to pay attention to information that confirms their belief and ignore information that contradicts it) could be used to develop a more holistic picture of the role of decision biases and cognitive limitations in early stage investing. Although we present results on a combination of all three biases, future research should also investigate different combinations of biases and their effect on early stage investment performance.

Sixth, although the angel investment platform we study is one of the largest in Europe and operates in many countries, most of the investments on this platform are made by European BAs in European companies. Because cultures differ in terms of investment practices, legislation, and risk-taking attitude, we encourage future studies to investigate angel investment platforms in other markets.

Finally, an important track of future research is how BAs can use ML to support their investment decisions. Examining investment cases in which BAs make collaborative decisions with a ML algorithm would increase our understanding of the potential of ML in early stage investing. Although research on hybrid intelligence suggests that combining the complementary capabilities of humans and machines enables superior predictions (e.g., Kamar, 2016), the positive and negative implications (e.g., preventing human learning processes) for entrepreneurship need further attention.

Conclusion

AI in general and ML algorithms in particular are entering many new decision environments, including entrepreneurship. The question of under which circumstances early stage investments would benefit from these new technologies is intriguing but remains unanswered to date. Our results show that our algorithm, on average, achieves higher investment performance when compared to the performance of BAs. In addition, we show that experienced BAs, especially those who use their experience to suppress their decision biases, can still produce best-in-class investment returns.

Appendix I: Model training

The general goal of our approach was to predict the survival of new ventures given a series of characteristics as predictor variables. We modeled survival as right-censored time-to-event data. We right-censored survival because we could only infer whether a new venture had died up to the end of our study period, that is, we did not know when surviving companies would potentially die in the future. In so doing, we measured the time of founding for each venture, and if a venture had died, we also inferred the time of death. If a venture was still alive, we considered the end of the study period as the censoring date. In our predictive model, we used a learning objective for the XGBoost algorithm that was based on a cox-proportional hazard model and a negative log loss function to deal with the right-censored time-to-event data (Hothorn et al., 2006). Generally, Cox proportional hazard functions estimate how a baseline hazard (the risk of dying) is influenced by a series of predictors at a given time t . Using the Cox proportional hazard learning objective, XGBoost returns all results on a hazard ratio scale that expresses how much the

baseline hazard across all startups under investigation at a given time t is scaled by a new venture's individual prediction (see Equation A1). Thus, our prediction approach is independent from any specific survival horizon (e.g., 5-year-survival).¹¹

$$h_t = h_0 \times \exp(\text{HR}) \quad (\text{A1})$$

where:

h_t = Hazard ratio of a given startup at time t

h_0 = Baseline hazard across all startups

HR = Marginal prediction for the given startup as weighted
sum of the individual decision tree predictions

In total, we had data on 623 startups; we had survival data for all of them, and for 51 startups, we also had data on their financial performance. We used the 51 startups as testing data and the remaining 572 startups as training data. In each bootstrap iteration, we randomly sampled 400 observations with replacement from the training data (ca. 70%; we verified that this is an appropriate size). In our training data, approximately 19% of companies were dead, so we made sure that each stratified sample contained 19% dead observations.

In each bootstrap iteration, we trained an individual XGBoost model using the Cox proportional hazard learning objective. For model training, we applied “early stopping” in order to prevent overfitting, i.e., 300 observations were used to train a model and 100 observations to determine when to stop training.¹² To train the XGBoost models, we used the following hyperparameters in all bootstrap iterations: 10 early stopping iterations, a learning rate of 0.0001, a maximal depth of 30, a minimum child weight of 0.001, a log loss reduction of 1, and a L1 regularization of 0. Following the general idea of ensembling, the predictions for the individual startups across all bootstrap iterations were aggregated by arithmetic mean. We used the R package “xgboost”.

Appendix 2: Text mining approach to LinkedIn profiles

To analyze the venture's human capital, we investigated the LinkedIn profiles of the founders applying a dictionary-based text mining approach, which is regularly used in entrepreneurship research (Benson et al., 2015; Prüfer & Prüfer, 2019; von Bloh et al., 2019). In this “bag-of-words” approach, a text is considered a collection of words and is compared against a dictionary (which may contain keywords associated with managerial experiences, thinking styles, emotions, or other quantities of interest) to determine how strongly a text is associated with concepts and constructs operationalized by the dictionary.

To analyze the textual LinkedIn data, we developed a total of 7 dictionaries (Table A1) reflecting the founders' prior professional experiences and education. We applied standard preprocessing steps to the textual data and dictionaries including tokenizing (determining individual positions of the LinkedIn profiles (e.g., a single job) and identifying individual words within these positions), removing stop words (e.g., words like “the” that occur in each text and do not have much meaning for comparing different texts), and stemming (reducing word inflections to their root form; e.g., “player” and “playing” will become “play”). In order to account for more complex word combinations and names with own meaning, we also created bi-grams (e.g., computer_science). Applying the dictionary-based bag-of-words analysis, we were able to determine how frequently the keywords from the dictionaries occurred in each position in each profile. Thus, we could determine whether the profiles suggested that the founders had certain professional experience or obtained a certain educational degree. For professional experiences, we

triangulated this information with the positions' reported length in months. If one position conveyed more than one type of experience, the reported length was split equally among the different types of experiences. All types of experiences were summed up for each founder. For the education variables, we created a set of dummy variables to capture whether a founder's position suggested that she or he had obtained an education (e.g., when a founder reported that she or he obtained a bachelor degree, we set the bachelor dummy to 1).

To develop the dictionaries, we applied (manual) content analysis to 100 randomly selected LinkedIn profiles in our sample. During this process, each position was manually classified based on whether it reflected a certain type of professional experience (and if so, how long that position was in months) or educational degree. In so doing, we created the same set of variables as used in our study in a manual fashion. These measures served as a benchmark in an iterative dictionary development process. We started by noting words that were frequently used to describe a certain type of experience or education (i.e., entrepreneurial experience was often described with the words "founder," "co-founder," "owner," etc.). We then applied the evolving dictionaries to the 100 profiles using our bag-of-words approach. In each iteration, we judged the reliability of the obtained measures by comparing them to the same measures that were determined in the manual coding procedure. After four iterations, we reached a Cronbach alpha of at least 0.7 for each measure, so we considered our measures to be reliable.

Acknowledgements

We thank the special issue editor Vincenzo Capizzi and three anonymous reviewers for their excellent guidance. We also thank Vivianna Fang He for her feedback on an earlier version of this manuscript.

Declaration of Conflicting Interests


The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

ORCID IDs

Ivo Blohm  <https://orcid.org/0000-0003-2422-5952>

Charlotta Sirén  <https://orcid.org/0000-0002-6984-4952>

Dietmar Grichnik  <https://orcid.org/0000-0001-7708-4773>

Notes

1. Our sample of BAs investing via our angel investment platform is comparable to samples used in other studies. For example, Bonini et al. (2018) reported a sample of platform members who were, on average, 48.32 years of age (our study: 48.77) and who had invested in, on average, 6.36 portfolio companies (our study: 4.92).
2. On average, the screening funnel looks as follows: Of 1,000 submitted investment proposals, around 400 are invited to provide additional information about their companies. Of those 400 ventures, 40 are invited to perform a live pitch to the BAs. After more detailed due diligence and negotiation with about 30% of the companies that pitch, around three companies receive an investment. Thus, the investment ratio is around 0.3%.
3. This notion refers to supervised ML (James et al., 2013), which we used in our study.

Table AI. Dictionaries for Analyzing LinkedIn Profiles

Variable	Word dictionary
Management experience	aufsichtsrat, board, cco, ceo, cfo, chairman, chairwoman, chief, cmo, cpo, cro, cto, director, executive, general_director, general_manager, geschaeftsfuehrer, head, lead, partner, projektleiter, president, verwaltungsrat, vorstand, vr
Entrepreneurial experience	businss_angel, cofounder, founder, founding_partner, gruender, inhaber, initiator, investor, mitgruender, owner, selbststaendiger_unternehmer, unternehmensinhaber
Technical experience	architect, business_intelligence, computer_science, cto, data, developer, development, engineer, engineering, programmer, systemadministrator, technical, technology
Consulting experience	accenture, bcg, boston_consulting, kearney, mck, mckinsey, oliver_wyman, roland_berger
Investment banking experience	barclays, citigroup, credit_suisse, deutsche_bank, goldman, hsbc, jp_morgan, merrill_lynnch, stanley, ubs
Bachelor's degree	b.a, b.ed, b.eng, b.sc, ba, bac.oec, bachelor, bachelor's, bs, bsc, ll.b, llb, dipl, diplom,
Master's degree	ll.m, llm, m.a, m.eng, m.f.a, m.mus, m.sc, ma, mag, magister, master, master's, meng, mlaw, msc,
MBA degree	cfa, emba, iema, imba, m.b.a, mba
PhD degree	doctor, doktor, dr, phd, ph.d.
Top university	bocconi, brown, cambridge, columbia, cornell, dartmouth, epfl, esade, eth, harvard, hec, hsg, ie_business, imd, imd, imperial_college, kit, lbs, london_business_school, oxford, pennsylvania, politecnico_di_milano, princeton, rotterdam_school, rwth, st_gallen, tu_berlin, tu_muenchen, tum, whu, yale
Technical education	bio, chemie, chemistry, computer, electronic, engineering, informatics, informatik, information, maschinenbau, material, mathematics, mechanical, physics, robotics, software, systemes_d'information, technology, wirtschaftsinformatik
Management education	accounting, banking, betriebssoekonomie, business, consulting, economia, economics, entrepreneurship, entreprise, finance, financial, management, marketing, sales, strategy, venture_capital, wirtschaftswissenschaften

4. Target leakage refers to biased predictive models that contain variables that have a direct association with the outcome, e.g., a model is trained on historic training data that contain variables that will not be available at the time of the prediction (Larsen & Becker, 2020).
5. Harrell's concordance index is frequently used to measure predictive accuracy of survival models. It ranks entities by their predicted survival times, that is, it compares the accuracy of predictions across all possible venture pairs in which at least one venture did not survive. This index ranges from 0 to 1, where 1 indicates complete agreement between predicted and actual survival times.
6. Values above 1,000 are rounded to the nearest integer.
7. Website: <https://archive.org/>, (last accessed May 8, 2020).
8. We used the Unicode Consortium's official emoji list: <https://unicode.org/emoji/charts/full-emoji-list.html> (last accessed February 21, 2020).
9. We performed a separate bootstrap analysis, in which we employed the same logic as in our main analysis but also included the ventures for which we had financial performance data into model

training. In 2,500 bootstrap iterations, we trained survival models on 400 ventures and again used survival predictions to build artificial portfolios based on 40 randomly selected ventures in the holdout data sets. However, we investigated the probability with which one or more ventures with known positive (negative) IRR were chosen by our ML approach. We compared these probabilities with the results of a “random investment process” that was modelled as hypergeometric distribution, that is, random portfolios of 3-15 ventures are drawn from a set of 40 ventures out of which one or more ventures have a known positive (negative) financial performance in each bootstrap iteration.

10. Our results should be interpreted in the light of the current stage of development of AI in general and ML algorithms in particular. As with every technology, the more these algorithms advance, the better tools they provide for decision making. However, even if these technologies become more capable in the future, our theoretical assumptions and implications persist: The potential of ML-based decision support is greatest for inexperienced and biased BAs, while experienced BAs who manage to suppress their cognitive biases remain the human benchmark in terms of early stage investment returns - also for more intelligent machines.
11. Website: <https://xgboost.readthedocs.io/en/latest/>, (last accessed May 8, 2020).
12. In early stopping, only part of the training data is actually used for training the model, while another part of the training data is used to determine when to stop model training (i.e., the validation data). As discussed in section 3.2., boosting is a sequential learning process. After each iteration in which the model is updated, the trained model is used to make predictions on the validation data. Observing the prediction performance on the validation set, the learning process is interrupted as soon as the prediction performance stops improving for a defined number of iterations and the model starts to overfit the data. Using this procedure, we made sure that ventures in the 400 sampled cases were either part of the training or the validation data.

References

- Agarwal, R., & Dhar, V. (2014). Editorial—big data, data science, and analytics: The opportunity and challenge for is research. *Information Systems Research*, 25(3), 443–448. <https://doi.org/10.1287/isre.2014.0546>
- Antretter, T., Blohm, I., & Grichnik, D. (2018). *Predicting startup survival from digital traces: Towards a procedure for early stage investors* [Conference session]. International Conference on Information Systems. San Francisco.
- Antretter, T., Blohm, I., Grichnik, D., & Wincent, J. (2019). Predicting new venture survival: A Twitter-based machine learning approach to measuring online legitimacy. *Journal of Business Venturing Insights*, 11(2019), e00109. <https://doi.org/10.1016/j.jbvi.2018.e00109>
- Baptista, R., Karaöz, M., & Mendonça, J. (2014). The impact of human capital on the early success of necessity versus opportunity-based entrepreneurs. *Small Business Economics*, 42(4), 831–847. <https://doi.org/10.1007/s11187-013-9502-z>
- Benjamin, G. A., & Margulis, J. (1996). *Finding your wings: How to locate private investors to fund your business*. Wiley.
- Benson, D. F., Brau, J. C., Cicon, J., & Ferris, S. P. (2015). Strategically camouflaged corporate governance in IPOs: Entrepreneurial masking and impression management. *Journal of Business Venturing*, 30(6), 839–864. <https://doi.org/10.1016/j.jbusvent.2015.03.001>
- Bessière, V., Stéphany, E., & Wirtz, P. (2020). Crowdfunding, business angels, and venture capital: An exploratory study of the concept of the funding trajectory. *Venture Capital*, 22(2), 135–160. <https://doi.org/10.1080/13691066.2019.1599188>
- Bickel, P. J., & Freedman, D. A. (1984). Asymptotic normality and the bootstrap in stratified sampling. *The Annals of Statistics*, 12(2), 470–482. <https://doi.org/10.1214/aos/1176346500>

- Blattberg, R. C., & Hoch, S. J. (1990). Database models and managerial intuition: 50% model + 50% manager. *Management Science*, *36*(8), 887–899. <https://doi.org/10.1287/mnsc.36.8.887>
- Böhm, M., Weking, J., Fortunat, F., Müller, S., Welp, I., & Krcmar, H. (2017). The business model DNA: Towards an approach for predicting business model success. *13th International Conference on Wirtschaftsinformatik* (pp. 1–15). St.Gallen. February 2017.
- Bonini, S., Capizzi, V., Valletta, M., & Zocchi, P. (2018). Angel network affiliation and business angels' investment practices. *Journal of Corporate Finance*, *50*(1), 592–608. <https://doi.org/10.1016/j.jcorpfin.2017.12.029>
- Boyer, T., & Blazy, R. (2014). Born to be alive? The survival of innovative and non-innovative French micro-start-ups. *Small Business Economics*, *42*(4), 669–683. <https://doi.org/10.1007/s11187-013-9522-8>
- Braga, A., & Logan, R. (2017). The emperor of strong AI has no clothes: Limits to artificial intelligence. *Information*, *8*(4), 156–177. <https://doi.org/10.3390/info8040156>
- Buchner, A., Mohamed, A., & Schwienbacher, A. (2017). Diversification, risk, and returns in venture capital. *Journal of Business Venturing*, *32*(5), 519–535. <https://doi.org/10.1016/j.jbusvent.2017.05.005>
- Capizzi, V. (2015). The returns of business angel investments and their major determinants. *Venture Capital*, *17*(4), 271–298. <https://doi.org/10.1080/13691066.2015.1092264>
- Carmona, P., Climent, F., & Momparler, A. (2019). Predicting failure in the US banking sector: An extreme gradient boosting approach. *International Review of Economics & Finance*, *61*(1), 304–323. <https://doi.org/10.1016/j.iref.2018.03.008>
- Carpentier, C., & Suret, J.-M. (2015). Angel group members' decision process and rejection criteria: A longitudinal analysis. *Journal of Business Venturing*, *30*(6), 808–821. <https://doi.org/10.1016/j.jbusvent.2015.04.002>
- Chang, Y.-C., Chang, K.-H., & Wu, G.-J. (2018). Application of eXtreme gradient boosting trees in the construction of credit risk assessment models for financial institutions. *Applied Soft Computing*, *73*(1), 914–920. <https://doi.org/10.1016/j.asoc.2018.09.029>
- Chen, G., Kim, K. A., Nofsinger, J. R., & Rui, O. M. (2007). Trading performance, disposition effect, overconfidence, representativeness bias, and experience of emerging market investors. *Journal of Behavioral Decision Making*, *20*(4), 425–451. <https://doi.org/10.1002/bdm.561>
- Chen, T., & Guestrin, C. (2016). Xgboost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, San Francisco. 2016. 785–794.
- Chi, M. T. H., Feltovich, P. J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, *5*(2), 121–152. https://doi.org/10.1207/s15516709cog0502_2
- Climent, F., Momparler, A., & Carmona, P. (2019). Anticipating bank distress in the Eurozone: An extreme gradient boosting approach. *Journal of Business Research*, *101*(1), 885–896. <https://doi.org/10.1016/j.jbusres.2018.11.015>
- Collewaert, V., & Manigart, S. (2016). Valuation of Angel-Backed companies: The role of Investor human capital. *Journal of Small Business Management*, *54*(1), 356–372.
- Cooper, A. C., Gimeno-Gascon, F. J., & Woo, C. Y. (1991). A resource-based prediction of new venture survival and growth. *Academy of Management Proceedings*, *1991*(1), 68–72. <https://doi.org/10.5465/ambpp.1991.4976561>
- Coval, J. D., & Moskowitz, T. J. (1999). Home bias at home: Local equity preference in domestic portfolios. *The Journal of Finance*, *54*(6), 2045–2073. <https://doi.org/10.1111/0022-1082.00181>
- Croce, A., Tenca, F., & Ughetto, E. (2017). How business angel groups work: Rejection criteria in investment evaluation. *International Small Business Journal*, *35*(4), 405–426. <https://doi.org/10.1177/0266242615622675>

- Cumming, D., & Dai, N. (2010). Local bias in venture capital investments. *Journal of Empirical Finance*, 17(3), 362–380. <https://doi.org/10.1016/j.jempfin.2009.11.001>
- Cumming, D., Fleming, G., & Schwienbacher, A. (2005). Liquidity risk and venture capital finance. *Financial Management*, 34(4), 77–105. <https://doi.org/10.1111/j.1755-053X.2005.tb00119.x>
- Delmar, F., & Shane, S. (2006). Does experience matter? The effect of founding team experience on the survival and sales of newly founded ventures. *Strategic Organization*, 4(3), 215–247. <https://doi.org/10.1177/1476127006066596>
- Dew, N., Read, S., Sarasvathy, S. D., & Wiltbank, R. (2009). Effectual versus predictive logics in entrepreneurial decision-making: Differences between experts and novices. *Journal of Business Venturing*, 24(4), 287–309. <https://doi.org/10.1016/j.jbusvent.2008.02.002>
- Dimov, D., & De Clercq, D. (2006). Venture capital investment strategy and portfolio failure rate: A longitudinal study. *Entrepreneurship Theory and Practice*, 30(2), 207–223. <https://doi.org/10.1111/j.1540-6520.2006.00118.x>
- Dimov, D., Shepherd, D. A., & Sutcliffe, K. M. (2007). Requisite expertise, firm reputation, and status in venture capital investment allocation decisions. *Journal of Business Venturing*, 22(4), 481–502. <https://doi.org/10.1016/j.jbusvent.2006.05.001>
- Dimov, D. P., & Shepherd, D. A. (2005). Human capital theory and venture capital firms: Exploring “home runs” and “strike outs”. *Journal of Business Venturing*, 20(1), 1–21. <https://doi.org/10.1016/j.jbusvent.2003.12.007>
- Drover, W., Busenitz, L., Matusik, S., Townsend, D., Anglin, A., & Dushnitsky, G. (2017). A review and road map of entrepreneurial equity financing research: Venture capital, corporate venture capital, angel investment, crowdfunding, and accelerators. *Journal of Management*, 43(6), 1820–1853. <https://doi.org/10.1177/0149206317690584>
- EBAN. (2017). *European early stage market statistics 2016*. European Business Angel Network.
- Efron, B., & Tibshirani, R. J. (1994). *An introduction to the bootstrap*. CRC Press Book.
- Einhorn, H. J. (1974). Expert judgment: Some necessary conditions and an example. *Journal of Applied Psychology*, 59(5), 562–571. <https://doi.org/10.1037/h0037164>
- Estrada, J. (2006). Downside risk in practice. *Journal of Applied Corporate Finance*, 18(1), 117–125. <https://doi.org/10.1111/j.1745-6622.2006.00080.x>
- Etter, M., Colleoni, E., Illia, L., Meggiorin, K., & D'Eugenio, A. (2018). Measuring organizational legitimacy in social media: Assessing citizens' judgments with sentiment analysis. *Business & Society*, 57(1), 60–97. <https://doi.org/10.1177/0007650316683926>
- Fischer, E., & Reuber, A. R. (2011). Social interaction via new social media: (How) can interactions on Twitter affect effectual thinking and behavior? *Journal of Business Venturing*, 26(1), 1–18. <https://doi.org/10.1016/j.jbusvent.2010.09.002>
- Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, 29(5), 1189–1232.
- Funk, R. J. (2014). Making the most of where you are: Geography, networks, and innovation in organizations. *Academy of Management Journal*, 57(1), 193–222. <https://doi.org/10.5465/amj.2012.0585>
- George, G., Osinga, E. C., Lavie, D., & Scott, B. A. (2016). Big data and data science methods for management research. *Academy of Management Journal*, 59(5), 1493–1507. <https://doi.org/10.5465/amj.2016.4005>
- Gervais, S., & Odean, T. (2001). Learning to be overconfident. *Review of Financial Studies*, 14(1), 1–27. <https://doi.org/10.1093/rfs/14.1.1>
- Giteno, J., Folta, T. B., Cooper, A. C., & Woo, C. Y. (1997). Survival of the fittest? Entrepreneurial human capital and the persistence of underperforming firms. *Administrative Science Quarterly*, 42(4), 750–783. <https://doi.org/10.2307/2393656>
- Harrell, F. E., Lee, K. L., & Mark, D. B. (1996). Multivariable prognostic models: Issues in developing models, evaluating assumptions and adequacy, and measuring and reducing errors. *Statistics in*

- Medicine*, 15(4), 361–387. [https://doi.org/10.1002/\(SICI\)1097-0258\(19960229\)15:4<361::AID-SIM168>3.0.CO;2-4](https://doi.org/10.1002/(SICI)1097-0258(19960229)15:4<361::AID-SIM168>3.0.CO;2-4)
- Harrison, R., Mason, C., & Robson, P. (2010). Determinants of long-distance investing by business angels in the UK. *Entrepreneurship & Regional Development*, 22(2), 113–137. <https://doi.org/10.1080/08985620802545928>
- Harrison, R. T., Mason, C., & Smith, D. (2015). Heuristics, learning and the business angel investment decision-making process. *Entrepreneurship & Regional Development*, 27(9-10), 527–554. <https://doi.org/10.1080/08985626.2015.1066875>
- Hastie, T., Tibshirani, R., & Friedman, J. (2009). Unsupervised learning. In T. Hastie, R. Tibshirani, & J. Friedman (Eds.), *The elements of statistical learning* (pp. 485–585). Springer.
- Hayward, M. L. A., Shepherd, D. A., & Griffin, D. (2006). A hubris theory of entrepreneurship. *Management Science*, 52(2), 160–172. <https://doi.org/10.1287/mnsc.1050.0483>
- Headd, B. (2003). Redefining business success: Distinguishing between closure and failure. *Small Business Economics*, 21(1), 51–61. <https://doi.org/10.1023/A:1024433630958>
- Heath, C., & Tversky, A. (1991). Preference and belief: Ambiguity and competence in choice under uncertainty. *Journal of Risk and Uncertainty*, 4(1), 5–28. <https://doi.org/10.1007/BF00057884>
- Hernandez, M., Raveendhran, R., Weingarten, E., & Barnett, M. (2019). How algorithms can diversify the startup pool. *MIT Sloan Management Review*, 61(1), 71–78.
- Hirshleifer, D. (2001). Investor psychology and asset pricing. *The Journal of Finance*, 56(4), 1533–1597. <https://doi.org/10.1111/0022-1082.00379>
- Hisrich, R. D., & Jankowicz, A. D. (1990). Intuition in venture capital decisions: An exploratory study using a new technique. *Journal of Business Venturing*, 5(1), 49–62. [https://doi.org/10.1016/0883-9026\(90\)90026-P](https://doi.org/10.1016/0883-9026(90)90026-P)
- Hochberg, Y. V., Ljungqvist, A., & Lu, Y. (2007). Whom you know matters: Venture capital networks and investment performance. *The Journal of Finance*, 62(1), 251–301. <https://doi.org/10.1111/j.1540-6261.2007.01207.x>
- Holcomb, T. R., Ireland, R. D., Holmes Jr., R. M., & Hitt, M. A. (2009). Architecture of entrepreneurial learning: Exploring the link among heuristics, knowledge, and action. *Entrepreneurship Theory and Practice*, 33(1), 167–192. <https://doi.org/10.1111/j.1540-6520.2008.00285.x>
- Hornuf, L., & Schmitt, M. (2016). Does a local bias exist in equity crowdfunding? *Max Planck Institute for Innovation & Competition Research Paper*, No. 16-07.
- Hothorn, T., Bühlmann, P., Dudoit, S., Molinaro, A., & van der Laan, M. J. (2006). Survival ensembles. *Biostatistics*, 7(3), 355–373. <https://doi.org/10.1093/biostatistics/kxj011>
- Hsu, D. H. (2004). What do entrepreneurs pay for venture capital affiliation? *The Journal of Finance*, 59(4), 1805–1844. <https://doi.org/10.1111/j.1540-6261.2004.00680.x>
- Huang, L. (2018). The role of investor gut feel in managing complexity and extreme risk. *Academy of Management Journal*, 61(5), 1821–1847. <https://doi.org/10.5465/amj.2016.1009>
- Huang, L., & Pearce, J. L. (2015). Managing the unknowable: The effectiveness of early-stage investor gut feel in entrepreneurial investment decisions. *Administrative Science Quarterly*, 60(4), 634–670.
- Huberman, G. (2001). Familiarity breeds investment. *Review of Financial Studies*, 14(3), 659–680. <https://doi.org/10.1093/rfs/14.3.659>
- Hunter, D., Saini, A., & Zaman, T. (2017). Picking winners: A framework for venture capital investment. *arXiv*, 1706.04229.
- Hyytinen, A., Pajarinen, M., & Rouvinen, P. (2015). Does innovativeness reduce startup survival rates? *Journal of Business Venturing*, 30(4), 564–581. <https://doi.org/10.1016/j.jbusvent.2014.10.001>
- IPEV. (2018). *International private equity and venture capital valuation guidelines*. British Private Equity & Venture Capital Association.

- Jacoby, J., Morrin, M., Johar, G., Gurhan, Z., Kuss, A., & Mazursky, D. (2001). Training novice investors to become more expert: The role of information accessing strategy. *Journal of Psychology and Financial Markets*, 2(2), 69–79.
- James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). *An introduction to statistical learning*. Springer.
- Jääskeläinen, M., & Maula, M. (2014). Do networks of financial intermediaries help reduce local bias? Evidence from Cross-border venture capital exits. *Journal of Business Venturing*, 29(5), 704–721. <https://doi.org/10.1016/j.jbusvent.2013.09.001>
- Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245), 255–260. <https://doi.org/10.1126/science.aaa8415>
- Kahneman, D., Knetsch, J. L., & Thaler, R. H. (1990). Experimental tests of the endowment effect and the coase theorem. *Journal of Political Economy*, 98(6), 1325–1348. <https://doi.org/10.1086/261737>
- Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, 80(4), 237–251. <https://doi.org/10.1037/h0034747>
- Kamar, E. (2016). Directions in hybrid intelligence: Complementing AI systems with human intelligence. *Proceedings of the International Joint Conference on Artificial Intelligence*, 1(2016), 4070–4073.
- Krishna, A., Agrawal, A., & Choudhary, A. (2016). *Predicting the outcome of startups: Less failure, more success* [Conference session]. International Conference on Data Mining Workshops, Barcelona.
- Kuppuswamy, V., & Bayus, B. L. (2017). Does my contribution to your crowdfunding project matter? *Journal of Business Venturing*, 32(1), 72–89. <https://doi.org/10.1016/j.jbusvent.2016.10.004>
- Larsen, K. R., & Becker, D. (2020). *Automated machine learning for business*. Oxford University Press.
- Li, J., & Tang, Y. (2010). CEO hubris and firm risk taking in China: The Moderating role of managerial discretion. *Academy of Management Journal*, 53(1), 45–68. <https://doi.org/10.5465/amj.2010.48036912>
- Li, T. R., Chamrajnagar, A. S., Fong, X. R., Rizik, N. R., & Fu, F. (2019). Sentiment-based prediction of alternative cryptocurrency price fluctuations using gradient boosting tree model. *Frontiers in Physics*, 7(7), 1–8. <https://doi.org/10.3389/fphy.2019.00098>
- Liew, J. K.-S., & Wang, G. Z. (2016). Twitter Sentiment and IPO performance: A cross-sectional examination. *The Journal of Portfolio Management*, 42(4), 129–135. <https://doi.org/10.3905/jpm.2016.42.4.129>
- Lindebaum, D., Vesa, M., & den Hond, F. (2020). Insights from the “machine stops” to better understand rational assumptions in algorithmic decision making and its implications for organizations. *Academy of Management Review*, 45(1), 247–263. <https://doi.org/10.5465/amr.2018.0181>
- Manigart, S., Baeyens, K., & Van Hyfte, W. (2002). The survival of venture capital backed companies. *Venture Capital*, 4(2), 103–124. <https://doi.org/10.1080/13691060110103233>
- Mason, C., & Harrison, R. (1994). The informal venture capital market in the UK. In A. Hughes & D. Storey (Eds.), *Financing small firms* (pp. 64–111). Routledge.
- Mason, C., & Harrison, R. (1996). Why 'business angels' say no: A case study of opportunities rejected by an informal investor syndicate. *International Small Business Journal: Researching Entrepreneurship*, 14(2), 35–51. <https://doi.org/10.1177/0266242696142003>
- Mason, C. M., & Harrison, R. T. (2002). Is it worth it? The rates of return from informal venture capital investments. *Journal of Business Venturing*, 17(3), 211–236. [https://doi.org/10.1016/S0883-9026\(00\)00060-4](https://doi.org/10.1016/S0883-9026(00)00060-4)
- Maxwell, A. L., Jeffrey, S. A., & Lévesque, M. (2011). Business angel early stage decision making. *Journal of Business Venturing*, 26(2), 212–225. <https://doi.org/10.1016/j.jbusvent.2009.09.002>
- McDuff, D., & Czerwinski, M. (2018). Designing emotionally sentient agents. *Communications of the ACM*, 61(12), 74–83. <https://doi.org/10.1145/3186591>
- Mitchell, C. J., De Houwer, J., & Lovibond, P. F. (2009). The propositional nature of human associative learning. *Behavioral and Brain Sciences*, 32(2), 183–198. <https://doi.org/10.1017/S0140525X09000855>

- Mittness, C. R., DeJordy, R., Ahuja, M. K., & Sudek, R. (2016). Extending the role of similarity attraction in friendship and advice networks in angel groups. *Entrepreneurship Theory and Practice*, 40(3), 627–655. <https://doi.org/10.1111/etap.12135>
- Moore, D. A., & Healy, P. J. (2008). The trouble with overconfidence. *Psychological Review*, 115(2), 502–517. <https://doi.org/10.1037/0033-295X.115.2.502>
- Navis, C., & Ozbek, O. V. (2016). The right people in the wrong places: The paradox of entrepreneurial entry and successful opportunity realization. *Academy of Management Review*, 41(1), 109–129. <https://doi.org/10.5465/amr.2013.0175>
- Newell, A., Shaw, J. C., & Simon, H. A. (1958). Elements of a theory of human problem solving. *Psychological Review*, 65(3), 151–166. <https://doi.org/10.1037/h0048495>
- Obschonka, M., & Audretsch, D. B. (2019). Artificial intelligence and big data in entrepreneurship: A new era has begun. *Small Business Economics*, 19(1), 1–11. <https://doi.org/10.1007/s11187-019-00202-4>
- Pan, S. J., & Yang, Q. (2009). A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10), 1345–1359. <https://doi.org/10.1109/TKDE.2009.191>
- Parhankangas, A., & Renko, M. (2017). Linguistic style and crowdfunding success among social and commercial entrepreneurs. *Journal of Business Venturing*, 32(2), 215–236. <https://doi.org/10.1016/j.jbusvent.2016.11.001>
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge University Press.
- Pennebaker, J. W., Francis, M. E., & Booth, R. J. (2001). *Linguistic inquiry and word count (LIWC): LIWC 2001*. Erlbaum.
- Pikulina, E., Renneboog, L., & Tobler, P. N. (2017). Overconfidence and investment: An experimental approach. *Journal of Corporate Finance*, 43(2), 175–192. <https://doi.org/10.1016/j.jcorpfin.2017.01.002>
- Plous, S. (1993). *The psychology of judgment and decision making*. McGraw-Hill Book Company.
- Pomeroy, J.-C. (1997). Artificial intelligence and human decision making. *European Journal of Operational Research*, 99(1), 3–25. [https://doi.org/10.1016/S0377-2217\(96\)00378-5](https://doi.org/10.1016/S0377-2217(96)00378-5)
- Prüfer, J., & Prüfer, P. (2019). Data science for entrepreneurship research: Studying demand dynamics for entrepreneurial skills in the Netherlands. *Small Business Economics*, (2019) 1–22.
- Raisch, S., & Krakowski, S. (2020). Artificial intelligence and management: The Automation-Augmentation paradox. *Academy of Management Review*, in press. <https://doi.org/10.5465/2018.0072>
- Raz, O., & Gloor, P. A. (2007). Size really matters—new insights for Start-ups’ survival. *Management Science*, 53(2), 169–177. <https://doi.org/10.1287/mnsc.1060.0609>
- Ruhnka, J. C., & Young, J. E. (1991). Some hypotheses about risk in venture capital investing. *Journal of Business Venturing*, 6(2), 115–133. [https://doi.org/10.1016/0883-9026\(91\)90014-5](https://doi.org/10.1016/0883-9026(91)90014-5)
- Sarasvathy, S. D. (2001). Causation and effectuation: Toward a theoretical shift from economic inevitability to entrepreneurial contingency. *Academy of Management Review*, 26(2), 243–263. <https://doi.org/10.5465/amr.2001.4378020>
- Schwab, A., & Zhang, Z. (2019). A new methodological frontier in entrepreneurship research: Big data studies. *Entrepreneurship Theory and Practice*, 43(5), 843–854. <https://doi.org/10.1177/1042258718760841>
- Scott, E. L., Shu, P., & Lubynsky, R. M. (2020). Entrepreneurial uncertainty and expert evaluation: An empirical analysis. *Management Science*, 66(3), 1278–1299. <https://doi.org/10.1287/mnsc.2018.3244>
- Shane, S. A. (2009). *Fool’s gold: The truth behind angel investing in America*. Oxford University Press.
- Shepherd, D. A., & Zacharakis, A. (2003). A new venture’s cognitive legitimacy: An assessment by customers. *Journal of Small Business Management*, 41(2), 148–167. <https://doi.org/10.1111/1540-627X.00073>
- Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1), 99–118. <https://doi.org/10.2307/1884852>

- Simon, D. P., & Simon, H. A. (1978). *Individual differences in solving physics problems*. Lawrence Erlbaum Associates.
- Sorenson, O., & Stuart, T. E. (2001). Syndication networks and the spatial distribution of venture capital investments. *American Journal of Sociology*, *106*(6), 1546–1588. <https://doi.org/10.1086/321301>
- Soto-Simeone, A., Sirén, C., & Antretter, T. (2020). New venture survival: A review and extension. *International Journal of Management Reviews*, *23*(1), 1–30. <https://doi.org/10.1111/ijmr.12229>
- Thaler, R. H., Tversky, A., Kahneman, D., & Schwartz, A. (1997). The effect of myopia and loss aversion on risk taking: An experimental test. *The Quarterly Journal of Economics*, *112*(2), 647–661. <https://doi.org/10.1162/003355397555226>
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*(4157), 1124–1131. <https://doi.org/10.1126/science.185.4157.1124>
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*(4), 297–323. <https://doi.org/10.1007/BF00122574>
- van Nieuwerburgh, S., & Veldkamp, L. (2009). Information immobility and the home bias puzzle. *The Journal of Finance*, *64*(3), 1187–1215. <https://doi.org/10.1111/j.1540-6261.2009.01462.x>
- van Praag, C. M. (2003). Business survival and success of young small business owners. *Small Business Economics*, *21*(1), 1–17. <https://doi.org/10.1023/A:1024453200297>
- van Witteloostuijn, A., & Kolkman, D. (2019). Is firm growth random? A machine learning perspective. *Journal of Business Venturing Insights*, *11*(1), e00107. <https://doi.org/10.1016/j.jbvi.2018.e00107>
- von Bloh, J., Broekel, T., Özgün, B., & Sternberg, R. (2019). New(s) data for entrepreneurship research? An innovative approach to use Big Data on media coverage. *Small Business Economics*, *25*(3), 1–22. <https://doi.org/10.1007/s11187-019-00209-x>
- Walthoff-Borm, X., Schwienbacher, A., & Vanacker, T. (2018). Equity crowdfunding: First resort or last resort? *Journal of Business Venturing*, *33*(4), 513–533. <https://doi.org/10.1016/j.jbusvent.2018.04.001>
- Westhead, P., & Wright, M. (1998). Novice, portfolio, and serial founders: Are they different? *Journal of Business Venturing*, *13*(3), 173–204. [https://doi.org/10.1016/S0883-9026\(97\)90002-1](https://doi.org/10.1016/S0883-9026(97)90002-1)
- Wetzel, W. E. (1981). Informal risk capital in New England. In K. H. Vesper (Ed.), *Frontiers of entrepreneurial research*. Babson College.
- Wigglesworth, R. (2016, January). Fintech: Search for a super-algo. *Financial Times*, 20.
- Wiltbank, R. (2005). Investment practices and outcomes of informal venture investors. *Venture Capital*, *7*(4), 343–357. <https://doi.org/10.1080/13691060500348876>
- Wiltbank, R., Read, S., Dew, N., & Sarasvathy, S. D. (2009). Prediction and control under uncertainty: Outcomes in angel investing. *Journal of Business Venturing*, *24*(2), 116–133. <https://doi.org/10.1016/j.jbusvent.2007.11.004>
- Witten, I. H., Frank, E., Hall, M. A., & Pal, C. J. (2016). *Data mining: Practical machine learning tools and techniques*. Morgan Kaufmann.
- Wolfe, M. T., & Shepherd, D. A. (2015). What do you have to say about that? Performance events and narratives' positive and negative emotional content. *Entrepreneurship Theory and Practice*, *39*(4), 895–925. <https://doi.org/10.1111/etap.12080>
- Wry, T., Lounsbury, M., & Jennings, P. D. (2014). Hybrid vigor: Securing venture capital by spanning categories in nanotechnology. *Academy of Management Journal*, *57*(5), 1309–1333. <https://doi.org/10.5465/amj.2011.0588>
- Yuan, F.-C., & Lee, C.-H. (2018). Investigating relationship between Google index and corporate profit using random forest. *Applied Mathematics*, *09*(1), 35–43. <https://doi.org/10.4236/am.2018.91004>
- Zacharakis, A. L., & Shepherd, D. A. (2001). The nature of information and overconfidence on venture capitalists' decision making. *Journal of Business Venturing*, *16*(4), 311–332. [https://doi.org/10.1016/S0883-9026\(99\)00052-X](https://doi.org/10.1016/S0883-9026(99)00052-X)

- Zarutskie, R. (2010). The role of top management team human capital in venture capital markets: Evidence from first-time funds. *Journal of Business Venturing*, 25(1), 155–172. <https://doi.org/10.1016/j.jbusvent.2008.05.008>
- Zimmerman, M. A., & Zeitz, G. J. (2002). Beyond survival: Achieving new venture growth by building legitimacy. *Academy of Management Review*, 27(3), 414–431. <https://doi.org/10.5465/amr.2002.7389921>

Author Biographies

Ivo Blohm is an Assistant Professor for Data Science and Management at the Institute for Information Management at the University of St. Gallen. His research focuses on business analytics, crowdsourcing and crowd work, as well as digital platforms. His work has appeared, among others, in such journals as *Information Systems Research*, *Journal of Management Information Systems*, and *California Management Review*.

Torben Antretter is a postdoctoral researcher at the University of St.Gallen. His research focusses on entrepreneurial finance, with particular focus on decision-making and early stage investment returns. His research is published in leading field journals such as the *Journal of Business Venturing* and the *International Journal of Management Reviews*. Above that, Torben is the founder of RightNow, a leading European consumer claims purchasing company.

Charlotta Sirén is an Associate Professor at University of St.Gallen. She studies entrepreneurial and strategic decision-making, with particular focus on the psychological and emotional aspects. Her research, which combines organizational and psychological theories has been published in internationally leading journals, such as *Journal of Management Studies*, *Journal of Business Venturing*, and *Entrepreneurship Theory and Practice*.

Dietmar Grichnik is Full Professor of Entrepreneurship and Director of the Institute of Technology Management at the University of St.Gallen. The main focus of his research is on entrepreneurial cognition and emotions, entrepreneurial finance, and entrepreneurial decision-making. His work has been published in leading journals such as *Journal of Management Studies*, *Journal of Business Venturing*, *Research Policy*, and *Entrepreneurship Theory and Practice*.

Joakim Wincent is a Professor of Entrepreneurship and Innovation at the Hanken School of Economics and at the University of St.Gallen. His current research focuses largely on strategic inter firm relations, corporate entrepreneurship and influences of psychological constructs such as passion and stress on innovative processes. His work is published in journals such as *Academy of Management Review*, *Strategic Management Journal*, *Journal of Management Studies* and *Organizational Research Methods*.